

Two Iterative Proximal-Point Algorithms for the Calculus of Divergence-based Estimators with Application to Mixture Models

Diaa Al Mohamad Michel Broniatowski,

A part of this work was presented in the conference paper [1]

Abstract

Estimators derived from an EM algorithm are not robust since they are based on the maximization of the likelihood function. We propose a proximal-point algorithm based on the EM algorithm which aim to minimize a divergence criterion. Resulting estimators are generally robust against outliers and misspecification. An EM-type proximal-point algorithm is also introduced in order to produce robust estimators for mixture models. Convergence properties of the two algorithms are treated. We relax an identifiability condition imposed on the proximal term in the literature; a condition which is generally not fulfilled by mixture models. The convergence of the introduced algorithms is discussed on a two-component Weibull mixture and a two-component Gaussian mixture entailing a condition on the initialization of the EM algorithm in order for the later to converge. Simulations on mixture models using different statistical divergences are provided to confirm the validity of our work and the robustness of the resulting estimators against outliers in comparison to the EM algorithm.

Index Terms

EM algorithm, mixture model, proximal-point algorithm, robustness, statistical divergence.

Diaa Al Mohamad is a PhD. student at Laboratoire de Statistique Théorique et Appliquée at the Univeristy of Paris 6 (UPMC) 4 place Jussieu 75005 Paris - France. Diaa Al Mohamad is the corresponding author of the article. email: diaa.almohamad@gmail.com

Michel Broniatowski is a Professor at Laboratoire de Statistique Théorique et Appliquée at the Univeristy of Paris 6 (UPMC) 4 place Jussieu 75005 Paris - France. Email: michel.broniatowski@upmc.fr

Two Iterative Proximal-Point Algorithms for the Calculus of Divergence-based Estimators with Application to Mixture Models

INTRODUCTION

The EM algorithm is a well known method for calculating the maximum likelihood estimator of a model where incomplete data are considered. For example, when working with mixture models in the context of clustering, the labels or classes of observations are unknown during the training phase. Several variants of the EM algorithm were proposed, see [2]. Another way to look at the EM algorithm is as a proximal point problem, see [3] and [4]. Indeed, one may rewrite the conditional expectation of the complete log-likelihood as a sum of the log-likelihood function and a distance-like function over the conditional densities of the labels provided an observation. Generally, the proximal term has a regularization effect in the sense that a proximal point algorithm is more stable and frequently outperforms classical optimization algorithms, see [5]. Chrétien and Hero [6] prove superlinear convergence of a proximal point algorithm derived by the EM algorithm. Notice that EM-type algorithms usually enjoy no more than linear convergence. Another aspect of proximal point algorithms is that they also permit avoiding saddle points as mentioned in [7].

Taking into consideration the need for robust estimators, and the fact that the MLE is the least robust estimator among the class of divergence-type estimators which we present below, we generalize the EM algorithm (and the version in [4]) by replacing the log-likelihood function by an estimator of a statistical divergence between the *true distribution* of the data and the model. We are particularly interested in φ -divergences and the density power divergence which is a Bregman divergence. We recall these two estimators briefly.

The density power divergence introduced by [8] is defined as follows:

$$D_a(g, f) = \int f^{1+a}(y) - \frac{a+1}{a}g(y)f^a(y) + \frac{1}{a}g^{1+a}(y)dy, \quad \text{with } a > 0, \quad (1)$$

for two probability density functions f and g . Given a random sample Y_1, \dots, Y_n distributed according to some probability measure P_T with density p_T with respect to the Lebesgue measure, and given a model $(p_\phi)_{\phi \in \Phi}$, the minimum density power divergence estimator (MDPD) is defined by:

$$\begin{aligned} \hat{\phi}_n &= \arg \inf_{\phi \in \Phi} \int p_\phi^{1+a}(z)dz - \frac{a+1}{a} \frac{1}{n} \sum_{i=1}^n p_\phi^a(Y_i) \\ &= \arg \inf_{\phi \in \Phi} \mathbb{E}_{P_\phi} [p_\phi^a] - \frac{a+1}{a} \mathbb{E}_{P_n} [p_\phi^a]. \end{aligned} \quad (2)$$

Consistency and robustness properties of the MDPD were studied by [8]. The authors show that, the MDPD is generally robust for $a > 0$ but the most interesting values of a are in the interval $(0, 1)$. Notice that when $a = 1$, the MDPD corresponds to the L^2 estimator, and as a goes to zero, we obtain the

MLE. See [9] for further properties.

A φ -divergence in the sense of Csiszár [10] is defined (see also [11]) by:

$$D_\varphi(Q, P) = \int \varphi \left(\frac{dQ}{dP}(y) \right) dP(y), \quad (3)$$

where φ is a nonnegative strictly convex function and Q and P are two probability measures such that Q is absolutely continuous with respect to P . Examples of such divergences are: the Kullback-Leibler (KL) divergence for $\varphi(t) = t \log(t) - t + 1$, the modified KL divergence for $\varphi(t) = -\log(t) + t - 1$, the hellinger distance for $\varphi(t) = \frac{1}{2}(\sqrt{t} - 1)$ among others. All these well-known divergences belong to the class of Cressie-Read functions defined by:

$$\varphi_\gamma(t) = \frac{x^\gamma - \gamma x + \gamma - 1}{\gamma(\gamma - 1)} \quad (4)$$

for $\gamma \in \mathbb{R} \setminus \{0, 1\}$ and $\varphi_1(t) = t \log(t) - t + 1$ and $\varphi_0(t) = -\log(t) + t - 1$.

Since the φ -divergence calculus uses the unknown true distribution, we need to estimate it. We consider the dual estimator of the divergence introduced independently by [12] and [13]. The use of this estimator is motivated by many reasons. Its minimum coincides with the MLE for $\varphi(t) = -\log(t) + t - 1$. Besides, it has the same form for discrete and continuous models, and does not consider any partitioning or smoothing which is not the case of other estimators such as [14], [15] and [16] which use kernel density estimators.

The dual estimator of the φ -divergence given an n -sample Y_1, \dots, Y_n is given by:

$$\hat{D}_\varphi(p_\phi, p_T) = \sup_{\alpha \in \Phi} \int \varphi' \left(\frac{p_\phi}{p_\alpha} \right) (x) p_\phi(x) dx - \frac{1}{n} \sum_{i=1}^n \varphi^\# \left(\frac{p_\phi}{p_\alpha} \right) (Y_i), \quad (5)$$

with $\varphi^\#(t) = t\varphi'(t) - \varphi(t)$. Al Mohamad [17] argues that this formula works well under the model, however, when we are not, this quantity largely underestimates the divergence between the true distribution and the model, and proposes following modification:

$$\tilde{D}_\varphi(p_\phi, p_T) = \int \varphi' \left(\frac{p_\phi}{K_{n,w}} \right) (x) p_\phi(x) dx - \frac{1}{n} \sum_{i=1}^n \varphi^\# \left(\frac{p_\phi}{K_{n,w}} \right) (Y_i), \quad (6)$$

where $K_{n,w}$ is a nonparametric estimator¹ of the true distribution P_T . The resulting new estimator is robust against outliers. It also permits to get rid of the supremal form which, as we will see later, entails technical and practical issues when one needs to use the continuity or the differentiability of $\hat{D}_\varphi(p_\phi, p_T)$ with respect to ϕ in order to prove the convergence of the algorithm.

The minimum dual φ -divergence estimator (MD φ DE) is defined as the argument of the infimum² of either $\hat{D}_\varphi(p_\phi, p_T)$ or $\tilde{D}_\varphi(p_\phi, p_T)$.

$$\text{Classical MD}\varphi\text{DE} = \arg \inf_{\phi \in \Phi} \hat{D}_\varphi(p_\phi, p_T), \quad (7)$$

$$\text{Kernel-based MD}\varphi\text{DE} = \arg \inf_{\phi \in \Phi} \tilde{D}_\varphi(p_\phi, p_T). \quad (8)$$

¹For example, and here in this paper, $K_{n,w}$ is a kernel density estimator based on either symmetric or asymmetric kernel (with or without bias-correction).

²Since there is no guarantee in general that the infimum is unique, the MD φ DE is defined as any of the points verifying the infimum.

Asymptotic properties and consistency of these two estimators can be found in [11] and [17]. Robustness properties were also studied using the influence function approach in [18] and [17]. The kernel-based MD φ DE (8) seems to be a *better* estimator than the classical MD φ DE (7) in the sense that the former is robust whereas the later is generally not. Under the model, the estimator given by (7) is, however, more efficient.

Here in this paper, we propose to calculate the two MD φ DEs and the MDPD using an iterative procedure based on the work of [4] on the log-likelihood function. This procedure has the form of a proximal point algorithm, and extends the EM algorithm. This algorithm was already introduced and discussed in [1] and [19]. We also propose in this paper a two-step iterative algorithm to calculate the MD φ DE for mixture models motivated by the EM algorithm. A step to calculate the proportion and a step to calculate the parameters of the components. Proofs for this simplified version become more technical. The goal of this simplification is to reduce the dimension over which we optimize since in lower dimensions, optimization procedures are more efficient³. Our convergence proof requires some regularity of the estimated divergence with respect to the parameter vector which is not simply checked using (5). Recent results in [20] provide sufficient conditions to solve this problem. Differentiability with respect to ϕ still remains a very hard task, therefore, our results cover cases when the objective function is not differentiable.

Another contribution of this work concerns the assumptions ensuring the convergence of the algorithm. In previous works on such type of proximal algorithms such as [4] and [3], the proximal term is supposed to verify an identifiability property. In other words $D(\phi, \phi') = 0$ if and only if $\phi = \phi'$. We show that such property is difficult and it is often not fulfilled in mixture models. We provide a way to relax such assumption in order to cover the case of distance-like functions such as the Kullback-Liebler (the EM case).

The paper is organized as follows: We explain in Section I the context and indicate the mathematical notations which may differ from standard ones. We also present the progression and the derivation of our set of algorithms from the EM algorithm and passing by Tseng's generalization. Section II is devoted to the analytical properties of a supremum function, i.e. a function defined as $\sup_{\alpha} f(\alpha, \phi)$ which is the case of the dual representation of the divergence presented above. In section III, we prove some convergence properties of the sequence generated by our algorithm. We show in Section IV and by examples, how one can prove convergence of the proposed algorithms in Gaussian and Weibull mixtures including a convergence proof of the EM algorithm. Finally, Section V gives some experimental results confirming the validity of the methods proposed in comparison simply to the maximum likelihood estimator calculated through the EM algorithm.

³This does not cover all optimization methods. For example, the Nelder-Mead algorithm is considered as "unreliable" in univariate optimization. The Brent method can be used as an alternative. Note that these two algorithms are suitable for not differentiable functions since they only use function values to reach an optimum.

I. A DESCRIPTION OF THE ALGORITHM

A. General context and notations

Let (X, Y) be a couple of random variables with joint probability density function $f(x, y|\phi)$ parametrized by a vector of parameters $\phi \in \Phi \subset \mathbb{R}^d$. Let $(X_1, Y_1), \dots, (X_n, Y_n)$ be n copies of (X, Y) independently and identically distributed. Finally, let $(x_1, y_1), \dots, (x_n, y_n)$ be n realizations of the n copies of (X, Y) . The x_i 's are the unobserved data (labels) and the y_i 's are the observations. The vector of parameters ϕ is unknown and need to be estimated.

The observed data y_i are supposed to be real vectors and the labels x_i belong to a space \mathcal{X} not necessarily finite unless mentioned otherwise. Denote dx the measure on the label space \mathcal{X} (for example the counting measure if \mathcal{X} is discrete). The marginal density of the observed data is given by $p_\phi(y) = \int f(x, y|\phi)dx$. For a parametrized function f with a parameter a , we write $f(x|a)$. We use the notation ϕ^k for sequences with the index above. Derivatives of a real valued function ψ defined on \mathbb{R} are written as ψ', ψ'' , etc. We use ∇f for the gradient of real function f defined on \mathbb{R}^d , ∂f to its subgradient and J_f to the matrix of second order partial derivatives. For a generic function H of two variables (ϕ, θ) , $\nabla_1 H(\phi, \theta)$ denotes the gradient with respect to the first (vectorial) variable ϕ .

B. EM algorithm and Tseng's generalization

The EM algorithm is a well-known method for calculating the maximum likelihood estimator of a model where incomplete data are considered. For example, when working with mixture models in the context of clustering, the labels or classes of observations are unknown during the training phase. Several variants of the EM algorithm were proposed, see [2]. The EM algorithm estimates the unknown parameter vector by generating the sequence (see [21]):

$$\begin{aligned}\phi^{k+1} &= \arg \max_{\Phi} Q(\phi, \phi^k) \\ &= \arg \max_{\Phi} \mathbb{E} [\log(f(\mathbf{X}, \mathbf{Y}|\phi)) \mid \mathbf{Y} = \mathbf{y}, \phi^k],\end{aligned}$$

where $\mathbf{X} = (X_1, \dots, X_n)$, $\mathbf{Y} = (Y_1, \dots, Y_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$. By independence between the couples (X_i, Y_i) 's, the previous iteration may be rewritten as:

$$\begin{aligned}\phi^{k+1} &= \arg \max_{\Phi} \sum_{i=1}^n \mathbb{E} [\log(f(X_i, Y_i|\phi)) \mid Y_i = y_i, \phi^k] \\ &= \arg \max_{\Phi} \sum_{i=1}^n \int_{\mathcal{X}} \log(f(x, y_i|\phi)) h_i(x|\phi^k) dx,\end{aligned}\tag{9}$$

where $h_i(x|\phi^k)$ is the conditional density of the labels (at step k) provided y_i . It is given by:

$$h_i(x|\phi^k) = \frac{f(x, y_i|\phi^k)}{p_{\phi^k}(y_i)}.\tag{10}$$

This justifies the recurrence equation given by [4]. It is slightly different from the EM recurrence defined in [21]. The conditional expectation of the logarithm of the complete likelihood provided the data and the parameter vector of the previous iteration is calculated, here, on the vector of observed data. The expectation is replaced by an integral against the corresponding conditional density of the labels.

It is well-known that the EM iterations can be rewritten as a difference between the log-likelihood and a *Kullback-Liebler* distance-like function. Indeed, using (10) in (9), one can write:

$$\begin{aligned}
\phi^{k+1} &= \arg \max_{\Phi} \sum_{i=1}^n \int_{\mathcal{X}} \log(h_i(x|\phi) \times p_{\phi}(y_i)) h_i(x|\phi^k) dx \\
&= \arg \max_{\Phi} \sum_{i=1}^n \int_{\mathcal{X}} \log(p_{\phi}(y_i)) h_i(x|\phi^k) dx + \sum_{i=1}^n \int_{\mathcal{X}} \log(h_i(x|\phi)) h_i(x|\phi^k) dx \\
&= \arg \max_{\Phi} \sum_{i=1}^n \log(p_{\phi}(y_i)) + \sum_{i=1}^n \int_{\mathcal{X}} \log\left(\frac{h_i(x|\phi)}{h_i(x|\phi^k)}\right) h_i(x|\phi^k) dx \\
&\quad + \sum_{i=1}^n \int_{\mathcal{X}} \log(h_i(x|\phi^k)) h_i(x|\phi^k) dx.
\end{aligned}$$

The final line is justified by the fact that $h_i(x|\phi)$ is a density, therefore it integrates to 1. The additional term does not depend on ϕ and, hence, can be omitted. We now have the following iterative procedure:

$$\phi^{k+1} = \arg \max_{\Phi} \sum_{i=1}^n \log(p_{\phi}(y_i)) + \sum_{i=1}^n \int_{\mathcal{X}} \log\left(\frac{h_i(x|\phi)}{h_i(x|\phi^k)}\right) h_i(x|\phi^k) dx. \quad (11)$$

As stated in [4], the previous iteration has the form of a proximal point maximization of the log-likelihood, i.e. a perturbation of the log-likelihood by a (modified) Kullback distance-like function defined on the conditional densities of the labels. Tseng proposed to generalize the Kullback distance-like term into other types of divergences. Tseng's recurrence is now defined by:

$$\phi^{k+1} = \arg \sup_{\phi} J(\phi) - D_{\psi}(\phi, \phi^k), \quad (12)$$

where J is the log-likelihood function and D_{ψ} is a distance-like function defined on the conditional probabilities of the classes provided the observations and is given by:

$$D_{\psi}(\phi, \phi^k) = \sum_{i=1}^n \int_{\mathcal{X}} \psi\left(\frac{h_i(x|\phi)}{h_i(x|\phi^k)}\right) h_i(x|\phi^k) dx, \quad (13)$$

for a real positive convex function ψ such that $\psi(1) = \psi'(1) = 0$. $D_{\psi}(\phi_1, \phi_2)$ is positive and equals zero if $\phi_1 = \phi_2$. Moreover, $D_{\psi}(\phi_1, \phi_2) = 0$ if and only if $\forall i, h_i(x|\phi_1) = h_i(x|\phi_2)$ dx -almost everywhere. Clearly, (12) and (11) are equivalent for $\psi(t) = -\log(t) + t - 1$.

C. Generalization of Tseng's algorithm

We use the relation between maximizing the log-likelihood and minimizing the Kullback-Liebler divergence to generalize the previous algorithm. We therefore replace the log-likelihood function by a divergence D between the true density of the data p_T and the model p_{ϕ} . This divergence will either denote density power divergence D_a given by (1) or a φ -divergence given by (3). Since the value of the divergence depends on the true density which is unknown, an estimator of the divergence needs to be considered. For the density power divergence, we use the optimized function in (2). For φ -divergences, we use the dual estimator of the divergence defined earlier by either (5) or (6). Denote \hat{D} for the estimator of the corresponding divergence. Our new algorithm is defined by the following recurrence:

$$\phi^{k+1} = \arg \inf_{\phi} \hat{D}(p_{\phi}, p_T) + \frac{1}{n} D_{\psi}(\phi, \phi^k) \quad (14)$$

where $D_\psi(\phi, \phi^k)$ is defined by (13). This algorithm was proposed in [19] in the context of φ -divergences. There is however no problem in defining the same algorithm for any statistical divergence family which generates the MLE. When $\varphi(t) = -\log(t) + t - 1$ or when $a \rightarrow 0$, it is easy to see that we get recurrence (12). Take for example the case of the approximation (5). Since $\varphi'(t) = \frac{-1}{t} + 1$, we have $\int \varphi' \left(\frac{p_\phi}{p_\alpha} \right) p_\phi dx = 0$. Hence,

$$\hat{D}_\varphi(p_\phi, p_T) = \sup_\alpha \frac{1}{n} \sum_{i=1}^n \log(p_\alpha(y_i)) - \frac{1}{n} \sum_{i=1}^n \log(p_\phi(y_i)).$$

Using the fact that the first term in $\hat{D}_\varphi(p_\phi, p_T)$ does not depend on ϕ , so it does not count in the $\arg \inf$ defining ϕ^{k+1} , we may rewrite (14) as:

$$\begin{aligned} \phi^{k+1} &= \arg \inf_{\phi} \left\{ \sup_{\alpha} \frac{1}{n} \sum_{i=1}^n \log(p_\alpha(y_i)) - \frac{1}{n} \sum_{i=1}^n \log(p_\phi(y_i)) + \frac{1}{n} D_\psi(\phi, \phi^k) \right\} \\ &= \arg \inf_{\phi} \left\{ -\frac{1}{n} \sum_{i=1}^n \log(p_\phi(y_i)) + \frac{1}{n} D_\psi(\phi, \phi^k) \right\} \\ &= \arg \sup_{\phi} \left\{ \frac{1}{n} \sum_{i=1}^n \log(p_\phi(y_i)) - \frac{1}{n} D_\psi(\phi, \phi^k) \right\} \\ &= \arg \sup_{\phi} J(\phi) - D_\psi(\phi, \phi^k). \end{aligned}$$

For notational simplicity, from now on, we redefine D_ψ with a normalization by n , i.e.

$$D_\psi(\phi, \phi^k) = \frac{1}{n} \sum_{i=1}^n \int_{\mathcal{X}} \psi \left(\frac{h_i(x|\phi)}{h_i(x|\phi^k)} \right) h_i(x|\phi^k) dx. \quad (15)$$

Hence, our set of algorithms is redefined by:

$$\phi^{k+1} = \arg \inf_{\phi} \hat{D}(p_\phi, p_T) + D_\psi(\phi, \phi^k). \quad (16)$$

We will see later that this iteration forces the estimated divergence to decrease and that under suitable conditions, it converges to a (local) minimum of $\hat{D}(p_\phi, p_T)$. It results that, algorithm (16) is a way to calculate the divergence-based estimators (2), (7) and (8).

Before proceeding to study the convergence properties of such algorithm, we will propose another algorithm for the case of mixture models. In the EM algorithm, the estimation of the parameters of a mixture model is done mainly by two steps. The first step estimates the proportions of the classes whereas the second step estimates the parameters defining the classes. Our idea is based on a directional optimization of the objective function in (16). Convergence properties of the two-step algorithm will also be studied, but the proofs are more technical.

D. A two-step EM-type algorithm for mixture models

Let p_ϕ be a mixture model with s components:

$$p_\phi(y) = \sum_{i=1}^s \lambda_i f_i(y|\theta_i). \quad (17)$$

Here, $\phi = (\lambda, \theta)$ with $\lambda = (\lambda_1, \dots, \lambda_s) \in [0, 1]^s$ such that $\sum_j \lambda_j = 1$, and $\theta = (\theta_1, \dots, \theta_s) \in \Theta \subset \mathbb{R}^{d-s}$ such that $\Phi \subset [0, 1]^s \times \Theta$. In the EM algorithm, the corresponding optimization to (16) can be solved by calculating an estimate of the λ 's as the proportions of classes, and then proceed to optimize on

the θ 's (see for example [22]). This simplifies the optimization in terms of complexity (optimization in lower spaces) and clarity (separate proportions from classes parameters). We want to build an algorithm with the same property and divide the optimization problem into two parts. One which estimates the proportions λ and another which estimates the parameters defining the form of each component θ . We propose the following algorithm:

$$\lambda^{k+1} = \arg \inf_{\lambda \in [0,1]^s, s.t. (\lambda, \theta^k) \in \Phi} \hat{D}(p_{\lambda, \theta^k}, p_T) + D_\psi((\lambda, \theta^k), \phi^k); \quad (18)$$

$$\theta^{k+1} = \arg \inf_{\theta \in \Theta, s.t. (\lambda^{k+1}, \theta) \in \Phi} \hat{D}(p_{\lambda^{k+1}, \theta}, p_T) + D_\psi((\lambda^{k+1}, \theta), \phi^k). \quad (19)$$

This algorithm corresponds to a directional optimization for recurrence (16). We can therefore prove analogously that the estimated divergence between the model and the true density decreases as we proceed with the recurrence.

We end the first part of this paper by three remarks:

- Function ψ defining the distance-like proximal term D_ψ needs not to be convex as in [4]. As we will see in the convergence proofs, the only properties needed are: ψ is a non negative function defined on \mathbb{R}_+ verifying $\psi(t) = 0$ iff $t = 1$, and $\psi'(t) = 0$ iff $t = 1$.
- The simplified version is not restricted to mixture models. Indeed, any parametric model, whose vector of parameters can be separated into two independent parts, can be estimated using the simplified version.
- As we will see in the proofs, results on the simplified version can be extended to a further simplified one. In other words, one may even consider an algorithm which attack a lower level of optimization. We may optimize on each class of the mixture model instead of the whole set of parameters. Since the analytic separation is not evident, one should expect some loss of quality as a cost of a less optimization time.

The remaining of the paper is devoted entirely to the study of the convergence of the sequences generated by either of the two sets of algorithms (16) and (18, 19) presented above. A key feature which will be needed in the proofs is the regularity of the objective function $\hat{D}_\varphi(p_\phi, p_T)$. This is the main goal of the following section.

II. ANALYTICAL DISCUSSION ABOUT THE REGULARITY OF THE ESTIMATED DIVERGENCE

The estimated divergence in (2) or (6) has an integral form. Thus, continuity and differentiability can be checked using Lebesgue theorems, see Section IV. However, the dual formula defining the estimator of the φ -divergence between the true density and the model (5) seems quite complicated. This is basically because of a functional integral and a supremum over it. Continuity and differentiation of the integral is resolved by Lebesgue theorems. We only need that the integrand as well as its partial derivatives to be uniformly bounded with respect to the parameter. However, continuity or differentiability of the supremum is more subtle. Indeed, even if the optimized function is \mathcal{C}^∞ , it does not imply the continuity

of its supremum. Take for example function $f(x, u) = -e^{xu}$. We have:

$$\sup_x f(x, u) = \begin{cases} -1 & \text{if } u = 0; \\ 0 & \text{if } u \neq 0. \end{cases}$$

On the basis of the theory presented in [20] about parametric optimization, we present two ways for studying continuity and differentiability of $\hat{D}_\varphi(p_\phi, p_{\phi_T})$ defined through (5). The first one is the most important because it is easier and demands less mathematical notations. In the first approach, we provide sufficient conditions in order to prove continuity and differentiability almost everywhere of the dual estimator of the divergence. This approach will be used in the study of the convergence of our proximal-point algorithm, see Section IV. The second approach is presented for the sake of completeness of the study. We give sufficient conditions which permit to prove the differentiability *everywhere*.

We recall first the definition of a subgradient of a real valued function f .

Definition 1 (Definition 8.3 in [20]). *Consider a function $f : \mathbb{R}^d \rightarrow \bar{\mathbb{R}}$ and a point ϕ^* with $f(\phi^*)$ finite. For a vector v in \mathbb{R}^d , one says that:*

(a) *v is a regular subgradient of f at ϕ^* , written $v \in \hat{\partial}f(\phi^*)$, if:*

$$f(\alpha) \geq f(\phi^*) + \langle v, \alpha - \phi^* \rangle + o(|\alpha - \phi^*|);$$

(b) *v is a (general) subgradient of f at ϕ^* , written $v \in \partial f(\phi^*)$, if there are sequences $\alpha^n \rightarrow \phi^*$ with $f(\alpha^n) \rightarrow f(\phi^*)$, and $v^n \in \hat{\partial}f(\alpha^n)$ with $v^n \rightarrow v$.*

A. A result of differentiability almost everywhere : Lower- \mathcal{C}^1 functions

Definition 2 ([20] Chap 10.). *A function $D : \Phi \rightarrow \mathbb{R}$, where Φ is an open set in \mathbb{R}^d , is said to be lower- \mathcal{C}^1 on Φ , if on some neighborhood V of each ϕ there is a representation*

$$D(\phi) = \sup_{\alpha \in T} f(\alpha, \phi)$$

in which the functions $\alpha \mapsto f(\alpha, \phi)$ are of class \mathcal{C}^1 on V and the set T is a compact set such that $f(\alpha, \phi)$ and $\nabla_\phi f(\alpha, \phi)$ depend continuously not just on $\phi \in \Phi$ but jointly on $(\alpha, \phi) \in T \times V$.

In our case, the supremum form is globally defined. Moreover, $T = \Phi$. In case Φ is bounded, it suffices then to take $T = cl(\Phi)$ the closure of Φ since $\alpha \mapsto f(\alpha, \phi)$ is continuous. The condition on T to be compact is essential here, and can not be compromised, so that it is necessary to reduce in a way or in another the optimization on α into a compact or at least a bounded set. For example, one may prove that the values of $\alpha \mapsto f(\alpha, \phi)$ near infinity are lower than some value inside Φ independently of ϕ .

Theorem 1 (Theorem 10.31 in [20]). *Any lower- \mathcal{C}^1 function D on an open set $\Phi \subset \mathbb{R}^d$ is both (strictly⁴)*

⁴A strictly continuous function f is a local Lipschitz continuous function, i.e. for each $x_0 \in \text{int}\Phi$, the following limit exists and is finite

$$\limsup_{x, x' \rightarrow x_0} \frac{|f(x') - f(x)|}{|x' - x|}$$

continuous and continuously differentiable where it is differentiable. Moreover, if Δ consists of the points where D is differentiable, then $\Phi \setminus \Delta$ is negligible⁵.

The stated result can be ensured by simple hypotheses on the model p_ϕ and the function φ . Unfortunately, since the estimated divergence $\hat{D}_\varphi(p_\phi, p_{\phi_T})$ will not be everywhere differentiable, we can no longer talk about the stationarity of $\hat{D}_\varphi(p_\phi, p_{\phi_T})$ at a limit point of the sequence ϕ^k generated for example by (16). We therefore, use the notion of subgradients. Indeed, when a function g is not differentiable, a necessary condition for x_0 to be a local minimum of g is that $0 \in \partial g(x_0)$ and it becomes sufficient whenever g is proper convex⁶. Moreover, as g becomes differentiable at x_0 , then $\nabla g(x_0) \in \partial g(x_0)$ with equality if and only if g is \mathcal{C}^1 . In other words, proving that $0 \in \partial \hat{D}_\varphi(p_{\hat{\phi}}, p_{\phi_T})$ means that $\hat{\phi}$ is a sort of a *generalized stationary point* of $\phi \mapsto \hat{D}_\varphi(p_\phi, p_{\phi_T})$.

We will be studying later on in paragraphs (IV-C1) and (V-A) examples where we verify with more details the previous conditions and see the resulting consequences on the sequence $(\phi^k)_k$.

B. A result of everywhere differentiability: Level-bounded functions

Definition 3 ([20] Chap 1.). *A function $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \bar{\mathbb{R}}$ with values $f(\alpha, \phi)$ is (upper) level-bounded in α locally uniformly in ϕ if for each ϕ_0 and $a \in \mathbb{R}$ there is a neighborhood V for ϕ_0 such that the set $\{(\alpha, \phi) | \phi \in V, f(\alpha, \phi) \geq a\}$ is bounded in $\mathbb{R}^d \times \mathbb{R}^d$ for every $a \in \mathbb{R}$.*

For a fixed ϕ , the level-boundedness property corresponds to having $f(\alpha, \phi) \rightarrow -\infty$ as $\|\alpha\| \rightarrow \infty$. In order to state the main result for this case, let ϕ_0 be a point at which we need to study continuity and differentiability of $\phi \mapsto \sup_\alpha f(\alpha, \phi)$. A first result gives sufficient conditions under which the supremum function is continuous. We state it as follows:

Theorem 2 ([20] Theorem 1.17). *Let $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \bar{\mathbb{R}}$ be an upper semicontinuous function. Suppose that $f(\alpha, \phi)$ is level-bounded in α locally uniformly in ϕ . For function $\phi \mapsto \sup_\alpha f(\alpha, \phi)$ to be continuous at ϕ_0 , a sufficient condition is the existence of $\alpha_0 \in \arg \max_\alpha f(\alpha, \phi_0)$ such that $\phi \mapsto f(\alpha_0, \phi)$ is continuous at ϕ_0 .*

Since in general, we do not know exactly where the supremum will be, one proves the continuity of $\phi \mapsto f(\alpha, \phi)$ for every α .

A Further result about continuity and differentiability of the supremum function can also be stated. Define, at first, the sets $Y(\phi_0)$ and $Y_\infty(\phi_0)$ as follows:

$$\begin{aligned} Y(\phi_0) &= \bigcup_{\alpha \in \arg \sup_\beta f(\beta, \phi_0)} M(\alpha, \phi_0), \quad \text{for } M(\alpha, \phi_0) = \{a | (0, a) \in \partial f(\alpha, \phi_0)\} \\ Y_\infty(\phi_0) &= \bigcup_{\alpha \in \arg \sup_\beta f(\beta, \phi_0)} M_\infty(\alpha, \phi_0), \quad \text{for } M_\infty(\alpha, \phi_0) = \{a | (0, a) \in \partial^\infty f(\alpha, \phi_0)\} \end{aligned}$$

⁵A set is called negligible if for every $\varepsilon > 0$, there is a family of boxes $\{B_k\}_k$ with d -dimensional volumes ε_k such that $A \subset \bigcup_k B_k$ and $\sum_k \varepsilon_k < \varepsilon$.

⁶See [20] theorem 10.1.

where $\partial^\infty f$ is the horizon subgradient, see Definition 8.3 (c) in [20]. We avoided to mention the definition here in order to keep the text clearer. Furthermore, in the whole chapter, the horizon subgradient will always be equal to the set $\{0\}$.

Theorem 3 (Corollary 10.14 in [20]). *For a proper upper semicontinuous function $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \bar{\mathbb{R}}$ such that $f(\alpha, \phi)$ is level-bounded in α locally uniformly in ϕ , and for $\phi_0 \in \text{dom } \sup_\alpha f(\alpha, \phi)$:*

- (a) *If $Y_\infty(\phi_0) = \{0\}$, then $\phi \mapsto \sup_\alpha f(\alpha, \phi)$ is strictly continuous at ϕ_0 ;*
- (b) *if $Y(\phi_0) = \{a\}$ too, then⁷ $\phi \mapsto \sup_\alpha f(\alpha, \phi)$ is \mathcal{C}^1 at ϕ_0 with $\nabla \sup_\alpha f(\alpha, \phi) = a$.*

In our examples, f will be a continuous function and even $\mathcal{C}^1(\Phi \times \Phi)$. This implies that $\partial^\infty f(\alpha, \phi) = \{0\}$ and $\partial f(\alpha, \phi) = \{\nabla f(\alpha, \phi)\}$, see Exercise 8.8 in [20]. Hence, $Y_\infty(\phi_0) = \{0\}$ whatever ϕ_0 in Φ . Moreover $M(\alpha, \phi_0) = \{\nabla_\phi f(\alpha, \phi_0)\}$ so that $Y(\phi_0) = \bigcup \{\nabla_\phi f(\alpha, \phi_0)\}$ and the union is on the set of suprema of $\alpha \mapsto f(\alpha, \phi_0)$. If $f(\alpha, \phi)$ is level-bounded in α locally uniformly in ϕ , then the supremum function becomes strictly continuous. Moreover, if the function f has the same gradient with respect to ϕ for all the suprema of $\alpha \mapsto f(\alpha, \phi)$, then $\sup_\alpha f(\alpha, \phi)$ becomes continuously differentiable. This is for example the case when function $\alpha \mapsto f(\alpha, \phi)$ has a unique global supremum for a fixed ϕ , which is for example the case of a strictly concave function (with respect to α for a fixed ϕ).

III. SOME CONVERGENCE PROPERTIES OF ϕ^k

We adapt the ideas given in [4] to develop a suitable proof for our proximal algorithm. We present some propositions which show how according to some possible situations one may prove convergence of the algorithms defined by recurrences (16) and (18, 19). Let $\phi^0 = (\lambda^0, \theta^0)$ be a given initialization for the parameters, and define the following set

$$\Phi^0 = \{\phi \in \Phi : \hat{D}(p_\phi, p_{\phi_T}) \leq \hat{D}(\phi^0, \phi_T)\}. \quad (20)$$

We suppose that Φ^0 is a subset of $\text{int}(\Phi)$. The idea of defining such set in this context is inherited from the paper of [23] which provided the first *correct proof* of convergence for the EM algorithm. Before going any further, we recall the following definition of a (generalized) stationary point.

Definition 4. *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a real valued function. If f is differentiable at a point ϕ^* such that $\nabla f(\phi^*) = 0$, we then say that ϕ^* is a stationary point of f . If f is not differentiable at ϕ^* but the subgradient of f at ϕ^* , say $\partial f(\phi^*)$, exists such that $0 \in \partial f(\phi^*)$, then ϕ^* is called a generalized stationary point of f .*

We will be using the following assumptions which will be checked in several examples later on.

A0. Functions $\phi \mapsto \hat{D}(p_\phi | p_{\phi_T}), D_\psi$ are lower semicontinuous;

⁷In the statement of the corollary in [20], the supremum function becomes strictly differentiable, but to avoid extra vocabularies, we replaced it with an equivalent property.

- A1. Functions $\phi \mapsto \hat{D}(p_\phi|p_{\phi_T})$, D_ψ and $\nabla_1 D_\psi$ are defined and continuous on, respectively, Φ , $\Phi \times \Phi$ and $\Phi \times \Phi$;
- AC. $\nabla \hat{D}(p_\phi|p_{\phi_T})$ is defined and continuous on Φ ;
- A2. Φ^0 is a compact subset of $\text{int}(\Phi)$;
- A3. $D_\psi(\phi, \bar{\phi}) > 0$ for all $\bar{\phi} \neq \phi \in \Phi$.

Recall also the assumptions on functions h_i defining D_ψ . We suppose that $h_i(x|\phi) > 0$, $dx - a.e.$, and $\psi(t) = 0$ iff $t = 1$. Besides $\psi'(t) = 0$ iff $t = 1$.

Concerning assumptions A1 and AC, we have previously discussed the analytical properties of $\hat{D}(p_\phi|p_{\phi_T})$ in Section II. In what concerns D_ψ , continuity and differentiability can be obtained merely by fulfilling Lebesgue theorems conditions. For example, if $h_i(x, \phi)$ is continuous and bounded uniformly away from 0 independently of ϕ , then continuity is guaranteed as soon as ψ is continuous. If we also suppose that $\nabla_\phi h_i(x, \phi)$ exists, is continuous and is uniformly bounded independently of ϕ , then as soon as ψ is continuously differentiable, D_ψ becomes continuously differentiable. For assumption A2, there is no universal method. Still, in all the examples that will be discussed later, we use the fact that the inverse image of a closed set by a continuous function is closed. Boundedness is usually ensured using a *suitable* choice of ϕ^0 . Finally, assumption A3 is checked using Lemma 2 proved in [4] which we restate here.

Lemma 1 (Lemma 2 in [4]). *Suppose ψ to be a continuous nonnegative function such that $\psi(t) = 0$ iff $t = 1$. For any ϕ and ϕ' in Φ , if $h_i(x|\phi) \neq h_i(x|\phi')$ for some $i \in \{1, \dots, n\}$ and some $x \in \text{int}(X)$ at which both $h_i(\cdot|\phi)$ and $h_i(\cdot|\phi')$ are continuous, then $D_\psi(\phi, \phi') > 0$.*

In section (IV), we present three different examples; a two-component Gaussian mixture, a two-component Weibull mixture and a Cauchy model. We will see that the Cauchy example verifies assumption A3. However, the Gaussian mixture does not seem to verify it. Indeed, the same fact stays true for any mixture of the exponential family.

We start by providing some general facts about the sequence $(\phi^k)_k$ and its existence. We also prove convergence of the sequence $(\hat{D}(p_{\phi^k}|p_{\phi_T}))_k$.

Remark 1. All results concerning algorithm (16) are proved even when assumption AC is not fulfilled. We give proofs using the subgradient of the estimated φ -divergence. In the case of the two-step algorithm (18, 19), it was not possible and thus remains an open problem. The difficulty resides in manipulating the *partial* subgradients with respect to λ and θ which cannot be handled in a similar way to the partial derivatives.

Remark 2. Convergence properties are proved without using the special form of the estimated φ -divergence. Thus, our theoretical approach applies to any optimization problem whose objective is to minimize a function $\phi \mapsto D(\phi)$. For example, our approach can be applied on density power divergences (2), (kernel-based) MD φ DE (7,8), Bregman divergences, S-divergences ([24]), Rényi pseudodistances (see for example [25]), etc.

The proofs of Propositions 1, 2 and 3 are only given for the two-step algorithm (18, 19). The proofs

of the case of algorithm (16) are direct adaptations of Theorem 1 and Lemme 1 in [4] for the case of the likelihood function, see also [19]. The proofs when assumption AC is not fulfilled can be found in [19] with \hat{D}_φ instead of \hat{D} .

Proposition 1. *We assume that recurrences (16) and (18, 19) are well defined in Φ . For both algorithms, the sequence $(\phi^k)_k$ verifies the following properties:*

- (a) $D_\varphi(p_{\phi^{k+1}}|p_T) \leq D_\varphi(p_{\phi^k}|p_T)$;
- (b) $\forall k, \phi^k \in \Phi^0$;
- (c) *Suppose that assumptions A0 and A2 are fulfilled, then the sequence $(\phi^k)_k$ is defined and bounded. Moreover, the sequence $(\hat{D}(\phi^k|\phi_T))_k$ converges.*

Proof: We prove (a). **For the two-step algorithm** defined by (18, 19), recurrence (18) and the definition of the arginf give:

$$\begin{aligned} \hat{D}(p_{\lambda^{k+1}, \theta^k}, p_T) + D_\psi((\lambda^{k+1}, \theta^k), \phi^k) &\leq \hat{D}(p_{\lambda^k, \theta^k}, p_T) + D_\psi((\lambda^k, \theta^k), \phi^k) \\ &\leq \hat{D}(p_{\lambda^k, \theta^k}, p_T). \end{aligned} \quad (21)$$

The second inequality is obtained using the fact that $D_\psi(\phi, \phi) = 0$. Using recurrence (19), we get:

$$\hat{D}(p_{\lambda^{k+1}, \theta^k}, p_T) + D_\psi((\lambda^{k+1}, \theta^k), \phi^k) \geq \hat{D}(p_{\lambda^{k+1}, \theta^{k+1}}, p_T) + D_\psi((\lambda^{k+1}, \theta^{k+1}), \phi^k) \quad (22)$$

$$\geq \hat{D}(p_{\lambda^{k+1}, \theta^{k+1}}, p_T). \quad (23)$$

The second inequality is obtained using the fact that $D(\phi|\phi') \geq 0$. The conclusion is reached by combining the two inequalities (21) and (23).

We prove (b). Using the decreasing property previously proved in (a), we have by recurrence $\forall k, \hat{D}(p_{\phi^{k+1}}, p_T) \leq \hat{D}(p_{\phi^k}, p_T) \leq \dots \leq \hat{D}(p_{\phi^0}, p_T)$. The result follows for both algorithms directly by definition of Φ^0 .

We prove (c). By induction on k . For $k = 0$, clearly $\phi^0 = (\lambda^0, \theta^0)$ is well defined (a choice we make⁸). Suppose for some $k \geq 0$ that $\phi^k = (\lambda^k, \theta^k)$ exists. **For the two-step algorithm** defined by (18, 19). The infimum in (18) can be calculated on λ 's such that $(\lambda, \theta^k) \in \Phi^0$. Indeed, suppose there exists a λ at which the value of the optimized function is less than its value at λ^k , i.e. $\hat{D}(p_{\lambda, \theta^k}, p_T) + D_\psi((\lambda, \theta^k), \phi^k) \leq \hat{D}(p_{\lambda^k, \theta^k}, p_T) + D_\psi((\lambda^k, \theta^k), \phi^k)$. We have:

$$\begin{aligned} \hat{D}(p_{\lambda, \theta^k}, p_T) &\leq \hat{D}(p_{\lambda, \theta^k}, p_T) + D_\psi((\lambda, \theta^k), \phi^k) \\ &\leq \hat{D}(p_{\lambda^k, \theta^k}, p_T) + D_\psi((\lambda^k, \theta^k), \phi^k) \\ &\leq \hat{D}(p_{\lambda^k, \theta^k}, p_T) \\ &\leq \hat{D}(p_{\phi^0}, p_T). \end{aligned}$$

This means that $(\lambda, \theta^k) \in \Phi^0$ and that the infimum needs not to be calculated for all values of λ in Φ , and can be restrained onto values which verify $(\lambda, \theta^k) \in \Phi^0$.

Define now $\Lambda_k = \{\lambda \in [0, 1]^s | (\lambda, \theta^k) \in \Phi^0\}$. First of all, $\lambda^k \in \Lambda_k$ since $(\lambda^k, \theta^k) \in \Phi^0$. Therefore, Λ_k is

⁸The choice of the initial point of the sequence may influence the convergence of the sequence. See the example of the Gaussian mixture in paragraph (IV-A).

not empty. Moreover, it is compact. Indeed, let $(\lambda^l)_l$ be a sequence of elements of Λ_k , then the sequence $((\lambda^l, \theta^k))_l$ is a sequence of elements of Φ^0 . By compactness of Φ^0 , there exists a subsequence which converges in Φ^0 to an element of the form $(\lambda^\infty, \theta^k)$ which clearly belongs to Λ_k . This proves that Λ_k is compact. Finally, since by assumption A0, the optimized function is lower semicontinuous so that it attains its infimum on the compact set Λ_k . We may now define λ^{k+1} as any vector verifying this infimum. The second part of the proof treats the definition of θ^{k+1} . Let θ be any vector such that $(\lambda^{k+1}, \theta) \in \Phi$ and at which the value of the optimized function in (19) is less than its value at ϕ^k . We have

$$\begin{aligned}
\hat{D}(p_{\lambda^{k+1}, \theta}, p_T) &\leq \hat{D}(p_{\lambda^{k+1}, \theta}, p_T) + D_\psi((\lambda^{k+1}, \theta), \phi^k) \\
&\leq \hat{D}(p_{\lambda^{k+1}, \theta^k}, p_T) + D_\psi((\lambda^{k+1}, \theta^k), \phi^k) \\
&\leq \hat{D}(p_{\lambda^k, \theta^k}, p_T) + D_\psi((\lambda^k, \theta^k), \phi^k) \\
&\leq \hat{D}(p_{\lambda^k, \theta^k}, p_T) \\
&\leq \hat{D}(p_{\phi^0}, p_T)
\end{aligned}$$

The third line comes from the previous definition of λ^{k+1} as an infimum of (18). This means that $(\lambda^{k+1}, \theta) \in \Phi^0$, and that the infimum in (19) can be calculated with respect to values θ which verifies $(\theta, \lambda^{k+1}) \in \Phi^0$. Define now $\Theta_k = \{\theta \in \mathbb{R}^{d-s} | (\lambda^{k+1}, \theta) \in \Phi^0\}$. One can prove analogously to Λ_k , that it is compact. The optimized function in (19) is, by assumption A0, lower semicontinuous so that its infimum is attained on the compact Θ_k . We may now define θ^{k+1} as any vector verifying this infimum.

Convergence of the sequence $(\hat{D}(p_{\phi^k}, p_T))_k$ in both algorithms comes from the fact that it is nonincreasing and bounded. It is nonincreasing by virtue of (a). Boundedness comes from the lower semicontinuity of $\phi \mapsto \hat{D}(p_\phi, p_T)$. Indeed, $\forall k, \hat{D}(p_{\phi^k}, p_T) \geq \inf_{\phi \in \Phi^0} \hat{D}(p_\phi, p_T)$. The infimum of a proper lower semicontinuous function on a compact set exists and is attained on this set. Hence, the quantity $\inf_{\phi \in \Phi^0} \hat{D}(p_\phi, p_T)$ exists and is finite. This ends the proof. ■

The interest of Proposition 1 is that the objective function is ensured, under mild assumptions, to decrease alongside the sequence $(\phi^k)_k$. This permits to build a stop criterion for the algorithm since in general there is no guarantee that the whole sequence $(\phi^k)_k$ converges. It may also continue to fluctuate in a neighborhood of an optimum. The following result provides a first characterization about the properties of the limit of the sequence $(\phi^k)_k$ as (generalized) a stationary point of the estimated φ -divergence.

Proposition 2. *Suppose that A1 is verified, and assume that Φ^0 is closed and $\{\phi^{k+1} - \phi^k\} \rightarrow 0$.*

- (a) *For both algorithms (16) and (18,19), if AC is verified, then the limit of every convergent subsequence is a stationary point of $\hat{D}(\cdot | p_T)$;*
- (b) *For the first algorithm (16), if $\hat{D}(\cdot | p_T)$ is not differentiable, then the limit of every convergent subsequence is a "generalized" stationary point of $\hat{D}(\cdot | p_T)$, i.e. zero belongs to the subgradient of $\hat{D}(\cdot | p_T)$ calculated at the limit point;*

Proof: We prove (a). Let $(\phi^{n_k})_k$ be a convergent subsequence of $(\phi^k)_k$ which converges to ϕ^∞ . First, $\phi^\infty \in \Phi^0$, because Φ^0 is closed and the subsequence (ϕ^{n_k}) is a sequence of elements of Φ^0 (proved in Proposition 1.b).

Let's show now that the subsequence (ϕ^{n_k+1}) also converges to ϕ^∞ . We simply have:

$$\|\phi^{n_k+1} - \phi^\infty\| \leq \|\phi^{n_k} - \phi^\infty\| + \|\phi^{n_k+1} - \phi^{n_k}\|$$

Since $\phi^{k+1} - \phi^k \rightarrow 0$ and $\phi^{n_k} \rightarrow \phi^\infty$, we conclude that $\phi^{n_k+1} \rightarrow \phi^\infty$.

For the two-step algorithm (18,19), by definition of λ^{n_k+1} and θ^{n_k+1} , they verify the infimum respectively in recurrences (18) and (19). Therefore, the gradient of the optimized function is zero for each step. In other words:

$$\begin{aligned} \nabla_\lambda \hat{D}(p_{\lambda^{n_k+1}, \theta^{n_k}}, p_T) + \nabla_\lambda D_\psi((\lambda^{n_k+1}, \theta^{n_k}), \phi^{n_k}) &= 0 \\ \nabla_\theta \hat{D}(p_{\lambda^{n_k+1}, \theta^{n_k+1}}, p_T) + \nabla_\theta D_\psi((\lambda^{n_k+1}, \theta^{n_k+1}), \phi^{n_k}) &= 0 \end{aligned}$$

Since both (ϕ^{n_k+1}) and (ϕ^{n_k}) converge to the same limit ϕ^∞ , then setting $\phi^\infty = (\lambda^\infty, \theta^\infty)$, we get λ^{n_k+1} and λ^{n_k} tends to λ^∞ . We also have θ^{n_k+1} and θ^{n_k} tends to θ^∞ . The continuity of the two gradients (assumptions A1 and AC) implies that:

$$\begin{aligned} \nabla_\lambda \hat{D}(p_{\lambda^\infty, \theta^\infty}, p_T) + \nabla_\lambda D_\psi((\lambda^\infty, \theta^\infty), \phi^\infty) &= 0 \\ \nabla_\theta \hat{D}(p_{\lambda^\infty, \theta^\infty}, p_T) + \nabla_\theta D_\psi((\lambda^\infty, \theta^\infty), \phi^\infty) &= 0 \end{aligned}$$

However, $\nabla D_\psi(\phi, \phi) = 0$, so that $\nabla_\lambda \hat{D}(p_{\phi^\infty}, p_T) = 0$ and $\nabla_\theta \hat{D}(p_{\phi^\infty}, p_T) = 0$. Hence $\nabla \hat{D}(p_{\phi^\infty}, p_T) = 0$.

We prove (b). See the proof of Proposition 2-b in [19]. ■

Proposition 3. *For both algorithms defined by (16) and (18,19), assume A1, A2 and A3 verified, then $\{\phi^{k+1} - \phi^k\} \rightarrow 0$. Thus, by Proposition 2 (according to whether AC is verified or not) implies that any limit point of the sequence ϕ^k is a (generalized)⁹ stationary point of $\hat{D}(\cdot | p_T)$.*

Proof: The arguments presented are the same for both algorithms (16) and (18,19). By contradiction, let's suppose that $\phi^{k+1} - \phi^k$ does not converge to 0. There exists a subsequence such that $\|\phi^{N_0(k)+1} - \phi^{N_0(k)}\| > \varepsilon$, $\forall k \geq k_0$. Since $(\phi^k)_k$ belongs to the compact set Φ^0 , there exists a convergent subsequence $(\phi^{N_1 \circ N_0(k)})_k$ such that $\phi^{N_1 \circ N_0(k)} \rightarrow \bar{\phi}$. The sequence $(\phi^{N_1 \circ N_0(k)+1})_k$ belongs to the compact set Φ^0 , therefore we can extract a further subsequence $(\phi^{N_2 \circ N_1 \circ N_0(k)+1})_k$ such that $\phi^{N_2 \circ N_1 \circ N_0(k)+1} \rightarrow \tilde{\phi}$. Besides $\bar{\phi} \neq \tilde{\phi}$. Finally since the sequence $(\phi^{N_1 \circ N_0(k)})_k$ is convergent, a further subsequence also converges to the same limit $\bar{\phi}$. We have proved the existence of a subsequence of $(\phi^k)_k$ such that $\phi^{N(k)+1} - \phi^{N(k)}$ does not converge to 0 and such that $\phi^{N(k)+1} \rightarrow \tilde{\phi}$, $\phi^{N(k)} \rightarrow \bar{\phi}$ with $\bar{\phi} \neq \tilde{\phi}$.

The real sequence $\hat{D}(p_{\phi^k}, p_T)_k$ converges as proved in Proposition 1-c. As a result, both sequences $\hat{D}(p_{\phi^{N(k)+1}}, p_T)$ and $\hat{D}(p_{\phi^{N(k)}}, p_T)$ converge to the same limit being subsequences of the same convergent sequence. In the proof of Proposition 1, we can deduce the following inequality:

$$\hat{D}(p_{\lambda^{k+1}, \theta^{k+1}}, p_T) + D_\psi((\lambda^{k+1}, \theta^{k+1}), \phi^k) \leq \hat{D}(p_{\lambda^k, \theta^k}, p_T) \quad (24)$$

which is also verified to any substitution of k by $N(k)$. By passing to the limit on k , we get $D_\psi(\tilde{\phi}, \bar{\phi}) \leq 0$. However, the distance-like function D_ψ is positive, so that it becomes zero. Using assumption A3,

⁹The case where AC is not verified is only proved for the first algorithm (16)

$D_\psi(\tilde{\phi}, \bar{\phi}) = 0$ implies that $\tilde{\phi} = \bar{\phi}$. This contradicts the hypothesis that $\phi^{k+1} - \phi^k$ does not converge to 0. The second part of the proposition is a direct result of Proposition 2. ■

We can go further in exploring the properties of the sequence $(\phi^k)_k$, but we need to impose more assumptions. The following corollary provides a convergence result of the *whole* sequence and not only some subsequence. The convergence is also towards a local minimum as soon as the estimated divergence is locally strictly convex.

Corollary 1. *Under assumptions of Proposition 3, the set of accumulation points of $(\phi^k)_k$ is a connected compact set. Moreover, if $\hat{D}(p_\phi, p_T)$ is strictly convex in a neighborhood of a limit point¹⁰ of the sequence $(\phi^k)_k$, then the whole sequence $(\phi^k)_k$ converges to a local minimum of $\hat{D}(p_\phi, p_T)$.*

Proof: The proof is based on Theorem 28.1 in [26], see [19]. ■

Proposition 3 although provides a general solution to prove that $\{\phi^{k+1} - \phi^k\} \rightarrow 0$, the identifiability assumption over the proximal term is hard to be fulfilled. It is not verified in the most simple mixtures such as a two component Gaussian mixture, see Section (IV-A).

This was the reason behind our next result. We prove that we do not need to assume identifiability of the proximal term in order to prove that any convergent subsequence of $(\phi^k)_k$ is a (generalized) stationary point of the estimated φ -divergence.

A similar idea was employed in [7] who studied a proximal algorithm for the log-likelihood function with a relaxation parameter¹¹. Their work however requires that the log-likelihood has $-\infty$ limit as $\|\phi\| \rightarrow \infty$ which is simply not verified on several mixture models (e.g. the Gaussian mixture model). Our result treat the problem from another approach based on the introduction of the set Φ^0 . The following result was already presented in the case of φ -divergences by [19], but since this result is still new, we prefer to rewrite the proof in the context of our paper.

Proposition 4. *Assume A1, AC and A2 verified. For the algorithm defined by (16), any convergent subsequence converges to a stationary point of the objective function $\phi \rightarrow \hat{D}(p_\phi, p_T)$. If AC is dropped, then 0 belongs to the subgradient of $\phi \mapsto \hat{D}(p_\phi, p_T)$ at the limit point.*

Proof: If $(\phi^k)_k$ converges to, say, ϕ^∞ , the result falls simply from Proposition 2.

If $(\phi^k)_k$ does not converge. Since Φ^0 is compact and $\forall k, \phi^k \in \Phi^0$ (proved in Proposition 1), there exists a subsequence $(\phi^{N_0(k)})_k$ such that $\phi^{N_0(k)} \rightarrow \tilde{\phi}$. Let's take the subsequence $(\phi^{N_0(k)-1})_k$. This subsequence does not necessarily converge; still it is contained in the compact Φ^0 , so that we can extract a further subsequence $(\phi^{N_1 \circ N_0(k)-1})_k$ which converges to, say, $\bar{\phi}$. Now, the subsequence $(\phi^{N_1 \circ N_0(k)})_k$ converges to $\tilde{\phi}$, because it is a subsequence of $(\phi^{N_0(k)})_k$. We have proved until now the existence of two convergent subsequences $\phi^{N(k)-1}$ and $\phi^{N(k)}$ with *a priori* different limits. For simplicity and without any loss of

¹⁰This assumption can be replaced by local strict convexity since *a priori*, we have no idea where might find a limit point of the sequence $(\phi^k)_k$.

¹¹A sequence of decreasing positive numbers multiplied by the proximal term.

generality, we will consider these subsequences to be ϕ^k and ϕ^{k+1} respectively.

Conserving previous notations, suppose that $\phi^{k+1} \rightarrow \tilde{\phi}$ and $\phi^k \rightarrow \bar{\phi}$. We use again inequality (24):

$$\hat{D}(p_{\phi^{k+1}}, p_T) + D_\psi(\phi^{k+1}, \phi^k) \leq \hat{D}(p_{\lambda^k, \theta^k}, p_T)$$

By taking the limits of the two parts of the inequality as k tends to infinity, and using the continuity of the two functions, we have

$$\hat{D}(p_{\tilde{\phi}}, p_T) + D_\psi(\tilde{\phi}, \bar{\phi}) \leq \hat{D}(p_{\bar{\phi}}, p_T)$$

Recall that under A1-2, the sequence $\left(\hat{D}(p_{\phi^k}, p_T)\right)_k$ converges, so that it has the same limit for any subsequence, i.e. $\hat{D}(p_{\tilde{\phi}}, p_T) = \hat{D}(p_{\bar{\phi}}, p_T)$. We also use the fact that the distance-like function D_ψ is nonnegative to deduce that $D_\psi(\tilde{\phi}, \bar{\phi}) = 0$. Looking closely at the definition of this divergence (15), we get that if the sum is zero, then each term is also zero since all terms are nonnegative. This means that:

$$\forall i \in \{1, \dots, n\}, \quad \int_{\mathcal{X}} \psi \left(\frac{h_i(x|\tilde{\phi})}{h_i(x|\bar{\phi})} \right) h_i(x|\bar{\phi}) dx = 0$$

The integrands are nonnegative functions, so they vanish almost everywhere with respect to the measure dx defined on the space of labels.

$$\forall i \in \{1, \dots, n\}, \quad \psi \left(\frac{h_i(x|\tilde{\phi})}{h_i(x|\bar{\phi})} \right) h_i(x|\bar{\phi}) = 0 \quad dx - a.e.$$

The conditional densities h_i are supposed to be positive¹², i.e. $h_i(x|\bar{\phi}) > 0, dx - a.e.$. Hence, $\psi \left(\frac{h_i(x|\tilde{\phi})}{h_i(x|\bar{\phi})} \right) = 0, dx - a.e.$. On the other hand, ψ is chosen in a way that $\psi(z) = 0$ iff $z = 1$, therefore :

$$\forall i \in \{1, \dots, n\}, \quad h_i(x|\tilde{\phi}) = h_i(x|\bar{\phi}) \quad dx - a.e. \quad (25)$$

Since ϕ^{k+1} is, by definition, an infimum of $\phi \mapsto \hat{D}(p_\phi, p_T) + D_\psi(\phi, \phi^k)$, then the gradient of this function is zero on ϕ^{k+1} . It results that:

$$\nabla \hat{D}(p_{\phi^{k+1}}, p_T) + \nabla D_\psi(\phi^{k+1}, \phi^k) = 0, \quad \forall k$$

Taking the limit on k , and using the continuity of the derivatives, we get that:

$$\nabla \hat{D}(p_{\tilde{\phi}}, p_T) + \nabla D_\psi(\tilde{\phi}, \bar{\phi}) = 0 \quad (26)$$

Let's write explicitly the gradient of the second divergence:

$$\nabla D_\psi(\tilde{\phi}, \bar{\phi}) = \sum_{i=1}^n \int_{\mathcal{X}} \frac{\nabla h_i(x|\tilde{\phi})}{h_i(x|\bar{\phi})} \psi' \left(\frac{h_i(x|\tilde{\phi})}{h_i(x|\bar{\phi})} \right) h_i(x|\bar{\phi})$$

We use now the identities (25), and the fact that $\psi'(1) = 0$, to deduce that:

$$\nabla D_\psi(\tilde{\phi}, \bar{\phi}) = 0$$

This entails using (26) that $\nabla \hat{D}(p_{\tilde{\phi}}, p_T) = 0$.

Comparing the proved result with the notation considered at the beginning of the proof, we have proved that the limit of the subsequence $(\phi^{N_1 \circ N_0(k)})_k$ is a stationary point of the objective function. Therefore, The final step is to deduce the same result on the original convergent subsequence $(\phi^{N_0(k)})_k$. This is

¹²In the case of two Gaussian (or more generally exponential) components, this is justified by virtue of a suitable choice of the initial condition.

simply due to the fact that $(\phi^{N_1 \circ N_0(k)})_k$ is a subsequence of the convergent sequence $(\phi^{N_0(k)})_k$, hence they have the same limit.

When assumption AC is dropped, the optimality condition in (16) implies :

$$-\nabla D_\psi(\phi^{k+1}, \phi^k) \in \partial \hat{D}(p_{\phi^{k+1}}, p_T) \quad \forall k$$

Function $\phi \mapsto \hat{D}(p_\phi, p_T)$ is continuous, hence its subgradient is outer semicontinuous and:

$$\limsup_{\phi^{k+1} \rightarrow \phi^\infty} \partial \hat{D}(p_{\phi^{k+1}}, p_T) \subset \partial \hat{D}(p_{\bar{\phi}}, p_T) \quad (27)$$

By definition of limsup:

$$\limsup_{\phi \rightarrow \phi^\infty} \partial \hat{D}(p_\phi, p_T) = \left\{ u \mid \exists \phi^k \rightarrow \phi^\infty, \exists u^k \rightarrow u \text{ with } u^k \in \partial \hat{D}(p_{\phi^k}, p_T) \right\}$$

In our scenario, $\phi = \phi^{k+1}$, $\phi^k = \phi^{k+1}$, $u = 0$ and $u^k = \nabla_1 D_\psi(\phi^{k+1}, \phi^k)$. We have proved above in this proof that $\nabla_1 D_\psi(\bar{\phi}, \bar{\phi}) = 0$ using only convergence of $(\hat{D}(p_{\phi^k}, p_T))_k$, inequality (24) and some properties of D_ψ . Assumption AC was not needed. Hence, $u^k \rightarrow 0$. This proves that, $u = 0 \in \limsup_{\phi^{k+1} \rightarrow \phi^\infty} \partial \hat{D}(p_{\phi^{k+1}}, p_T)$. Finally, using the inclusion (27), we get our result:

$$0 \in \partial \hat{D}(p_{\bar{\phi}}, p_T)$$

■

We could not perform the same idea on the two-step algorithm (18,19) without assuming that the difference between two consecutive terms of either the sequence of weights $(\lambda^k)_k$ or the sequence of form parameters $(\theta^k)_k$ converges to zero. Besides, when assumption AC is dropped, the proof becomes very complicated because we are obliged to work with partial subgradients. The problem is that the subgradient is a set-valued function and if zero belongs to both the partial subgradients with respect to λ and θ of the objective function, there is no guarantee that it belongs to the "whole" subgradient of the objective function. Hence, we do not have the elements of proof for such result for the time being.

Proposition 5. *Assume A1 and A2 verified. For the algorithm defined by (18,19). If $\|\theta^{k+1} - \theta^k\| \rightarrow 0$, then any convergent subsequence $(\phi^{N(k)})_k$ converges to a stationary point of the objective function $\phi \rightarrow \hat{D}(p_\phi, p_T)$.*

Proof: We prove (a). We use the same lines from the previous proof to deduce the existence of two convergent subsequences $\phi^{N(k)-1}$ and $\phi^{N(k)}$ with *a priori* different limits. For simplicity and without any loss of generality, we will consider these subsequences to be ϕ^k and ϕ^{k+1} respectively. Suppose that $\phi^k \rightarrow \bar{\phi} = (\bar{\lambda}, \bar{\theta})$ and $\phi^{k+1} \rightarrow \tilde{\phi} = (\tilde{\lambda}, \tilde{\theta})$.

We first use inequality (24) as in the previous proposition, the convergence of the sequence $(\hat{D}(p_{\lambda^k, \theta^k}, p_T))_k$ and some basic properties of D_ψ to deduce that:

$$\forall i \in \{1, \dots, n\}, \quad h_i(x|\tilde{\phi}) = h_i(x|\bar{\phi}) \quad dx - a.e. \quad (28)$$

Let's calculate the gradient of the objective function with respect to λ and θ separately at the limit of $(\phi^{k+1})_k$. By definition of θ^{k+1} as an arginf in (19), we have:

$$\frac{\partial}{\partial \theta} \hat{D}(p_{\lambda^{k+1}, \theta^{k+1}}, p_T) + \frac{\partial}{\partial \theta} D_\psi((\lambda^{k+1}, \theta^{k+1}), \phi^k) = 0 \quad \forall k$$

Using the continuity of the derivatives (Assumptions A1 and AC), we may pass to the limit inside the gradients:

$$\frac{\partial}{\partial \theta} \hat{D}(p_{\tilde{\lambda}, \tilde{\theta}}, p_T) + \frac{\partial}{\partial \theta} D_\psi((\tilde{\lambda}, \tilde{\theta}), \bar{\phi}) = 0 \quad \forall k$$

As in the proof of Proposition 3, all terms in the gradient of D_ψ depend on $\psi' \left(\frac{h_i(x|\tilde{\lambda}, \tilde{\theta})}{h_i(x|\bar{\phi})} \right)$ which is zero by virtue of (28). Hence $\frac{\partial}{\partial \theta} \hat{D}(p_{\tilde{\lambda}, \tilde{\theta}}, p_T) = 0$.

We prove now that $\frac{\partial}{\partial \lambda} \hat{D}(p_{\tilde{\lambda}, \tilde{\theta}}, p_T) = 0$. This is basically ensured by recurrence (18), identities (28), assumptions A1-AC and the fact that $\psi'(1) = 0$. Indeed, using recurrence (18), λ^{k+1} is an optimum so that the gradient of the objective function is zero:

$$\frac{\partial}{\partial \lambda} \hat{D}(p_{\lambda^{k+1}, \theta^k}, p_T) + \frac{\partial}{\partial \lambda} D_\psi((\lambda^{k+1}, \theta^k), \lambda^k, \theta^k) = 0, \quad \forall k$$

Since $\|\lambda^{k+1} - \theta^k\| \rightarrow 0$, then $\bar{\theta} = \tilde{\theta}$. By passing to the limit in the previous identity and using the continuity of the derivatives, we have:

$$\frac{\partial}{\partial \lambda} \hat{D}(p_{\tilde{\lambda}, \tilde{\theta}}, p_T) + \frac{\partial}{\partial \lambda} D_\psi((\tilde{\lambda}, \tilde{\theta}), \bar{\lambda}, \bar{\theta}) = 0$$

Since the derivative of D_ψ is a sum of terms which depend all on $\psi' \left(\frac{h_i(x|\tilde{\lambda}, \tilde{\theta})}{h_i(\bar{\lambda}, \bar{\theta})} \right)$, and using identities (28), we conclude that $\psi' \left(\frac{h_i(\bar{\lambda}, \bar{\theta})}{h_i(\bar{\lambda}, \bar{\theta})} \right) = \psi'(1) = 0$ and $\frac{\partial}{\partial \lambda} D_\psi((\tilde{\lambda}, \tilde{\theta}), \bar{\lambda}, \bar{\theta}) = 0$. Finally, $\bar{\theta} = \tilde{\theta}$ implies that $\frac{\partial}{\partial \lambda} \hat{D}(p_{\tilde{\lambda}, \tilde{\theta}}, p_T) = 0$.

We have proved that $\frac{\partial}{\partial \lambda} \hat{D}(p_{\tilde{\lambda}, \tilde{\theta}}, p_T) = 0$ and $\frac{\partial}{\partial \theta} \hat{D}(p_{\tilde{\lambda}, \tilde{\theta}}, p_T) = 0$, so the gradient is zero and the stated result is proved. \blacksquare

Remark 3. The previous proposition demands a condition on the distance between two consecutive members of the sequence $(\theta^k)_k$ which is *a priori* weaker than the same condition on the whole sequence $\phi^k = (\lambda^k, \theta^k)$. Still, as the regularization term D_ψ does not verify the identifiability condition A3, it stays an open problem for a further work. It is interesting to notice that condition $\|\theta^{k+1} - \theta^k\| \rightarrow 0$ can be replaced by $\|\lambda^{k+1} - \lambda^k\| \rightarrow 0$, but we then need to change the order of steps (18) and (19). A condition over the proportions seems to be *simpler*.

Remark 4. We can define an algorithm which converges to a global infimum of the estimated φ -divergence. The idea is very simple. We need to multiply the proximal term by a sequence $(\beta_k)_k$ of positive numbers which decreases to zero, for example $\beta_k = 1/k$. The justification of such variant can be deduced from Theorem 3.2.4 in [7]. The problem with this approach is that it depends heavily on the fact that the supremum on each step of the algorithm is calculated exactly. This does not happen in general unless function $\hat{D}(p_\phi, p_T) + \beta_k D_\psi(\phi, \phi^k)$ is strictly convex. Although in our approach, we use similar assumption to prove the consecutive decreasing of $\hat{D}(p_\phi, p_T)$, we can replace the infimum calculus in (16) by two things. We require at each step that we find a local infimum of $\hat{D}(p_\phi, p_T) + D_\psi(\phi, \phi^k)$ whose evaluation with $\phi \mapsto \hat{D}(p_\phi, p_T)$ is less than the previous term of the sequence ϕ^k . If we can no longer find any local maxima verifying the claim, the procedure stops with $\phi^{k+1} = \phi^k$. This ensures the availability of all proofs presented in this paper with no further changes.

IV. EXAMPLES

A. Two-component Gaussian mixture

We suppose that the model $(p_\phi)_{\phi \in \Phi}$ is a mixture of two Gaussian densities, and suppose that we are only interested in estimating the means $\mu = (\mu_1, \mu_2) \in \mathbb{R}^2$ and the proportions $\lambda = (\lambda_1, \lambda_2) \in [\eta, 1 - \eta]^2$. The use of η is to avoid cancellation of any of the two components and to keep the hypothesis about the conditional densities h_i true, i.e. $h_i(x|\phi) > 0$ for $x = 1, 2$. We also suppose to simplify the calculus that the components variances are reduced ($\sigma_i = 1$). The model takes the form:

$$p_{\lambda, \mu}(x) = \frac{\lambda}{\sqrt{2\pi}} e^{-\frac{1}{2}(x-\mu_1)^2} + \frac{1-\lambda}{\sqrt{2\pi}} e^{-\frac{1}{2}(x-\mu_2)^2}, \quad (29)$$

where $\Phi = [\eta, 1 - \eta]^s \times \mathbb{R}^s$. Here $\phi = (\lambda, \mu_1, \mu_2)$. The distance-like function D_ψ is defined by:

$$D_\psi(\phi, \phi^k) = \sum_{i=1}^n \psi \left(\frac{h_i(1|\phi)}{h_i(1|\phi^k)} \right) h_i(1|\phi^k) + \sum_{i=1}^n \psi \left(\frac{h_i(2|\phi)}{h_i(2|\phi^k)} \right) h_i(2|\phi^k),$$

where:

$$h_i(1|\phi) = \frac{\lambda e^{-\frac{1}{2}(y_i - \mu_1)^2}}{\lambda e^{-\frac{1}{2}(y_i - \mu_1)^2} + (1 - \lambda) e^{-\frac{1}{2}(y_i - \mu_2)^2}}, \quad h_i(2|\phi) = 1 - h_i(1|\phi).$$

It is clear that functions h_i are of class \mathcal{C}^1 on $(\text{int}(\Phi))$, and as a consequence, D_ψ is also of class \mathcal{C}^1 on $(\text{int}(\Phi))$.

If we use the MDPD (2), then function $\phi \mapsto \hat{D}(p_\phi, p_T)$ is clearly continuously differentiable by Lebesgue theorems. Recall that $\hat{D}_a(p_\phi, p_T)$ is given by $\int p_{\lambda, \mu}^{1+a}(y) dy - \frac{a+1}{a} \frac{1}{n} \sum p_{\lambda, \mu}^a(y_i)$, since we dropped the supplementary term $\frac{1}{a} \int p_T^{1+a}(y) dy$ from (1) because it does not depend on the parameters. Notice that for any $\mu = (\mu_1, \mu_2)$ such that $\|\mu\| < M$, $p_\phi^{1+a}(y) \leq c e^{-(1+a)y^2}$ for some positive constant c which depends on M and a . Thus, assumptions A1 and AC are verified. In order to prove that Φ^0 is compact, we prove that it is closed and bounded in the complete space $[\eta, 1 - \eta] \times \mathbb{R}^2$. Closedness is an immediate result of the continuity of the estimated divergence. Indeed,

$$\begin{aligned} \Phi^0 &= \left\{ \phi \in \Phi, \hat{D}_a(p_\phi, p_T) \leq \hat{D}_a(p_{\phi^0}, p_T) \right\} \\ &= \hat{D}_a(p_\phi, p_T)^{-1} \left((-\infty, \hat{D}_a(p_{\phi^0}, p_T)] \right). \end{aligned}$$

In order to ensure boundedness of Φ^0 , we need to choose carefully the initial point (λ^0, μ^0) of the algorithm. Since λ is bounded by 0 and 1, we only need to verify the boundedness of the means. If both means μ_1 and μ_2 go to $\pm\infty$, then $\hat{D}_a(p_{\phi, p_T}) \rightarrow 0$. Besides, if either of the means go to $\pm\infty$, then the corresponding component vanishes. Thus if we choose (λ^0, μ^0) such that:

$$\hat{D}_a(p_{(\lambda^0, \mu_1^0, \mu_2^0)}, p_T) < \min \left(0, \inf_{\lambda \in [\eta, 1-\eta], \mu_1 \in \mathbb{R}} \hat{D}_a(p_{(\lambda, \mu_1, \infty)}, p_T) \right), \quad (30)$$

then by definition of Φ^0 , any point of it must have a corresponding value of $\hat{D}_a(p_{\phi, p_T})$ less than its values at the extremities, i.e. when either of both means goes to infinity. Thus, under condition (30), Φ^0 is bounded. Now that assumption A2 is also fulfilled, we arrive to the following conclusion.

Conclusion 1. *Using Propositions 1 and 4 and under condition (30), the sequence $(\hat{D}_a(p_{\phi^k}, p_T))_k$ converges and there exists a subsequence $(\phi^{N(k)})$ which converges to a stationary point of the estimated*

divergence. Moreover, every limit point of the sequence $(\phi^k)_k$ is a stationary point of the estimated divergence.

If we are using the dual estimator of the φ -divergence given by (5). This was discussed in [19]. We cite only the final conclusion for the sequence $(\phi^k)_k$ defined by any of the proximal-point algorithms (16) or (18,19).

Conclusion 2. Using Propositions 1 and 4, if $\Phi = [\eta, 1 - \eta] \times [\mu_{\min}, \mu_{\max}]^2$, the sequence $(\hat{D}_\varphi(p_{\phi^k}, p_T))_k$ defined through formula (5) converges and there exists a subsequence $(\phi^{N(k)})$ which converges to a stationary point of the estimated divergence. Moreover, every limit point of the sequence $(\phi^k)_k$ is a stationary point of the estimated divergence.

If we are using the kernel-based dual estimator given by (6) with a Gaussian kernel density estimator, then if the we initialize any of the proximal-point algorithms (16) or (18,19) with ϕ^0 verifying:

$$\hat{D}_\varphi(p_\phi, p_T) < \min \left(\frac{1}{\gamma(\gamma - 1)}, \inf_{\lambda, \mu} \hat{D}_\varphi(p_{(\lambda, \infty, \mu)}, p_T) \right) \quad \text{if } \gamma \in (0, \infty) \setminus \{1\}; \quad (31)$$

$$\hat{D}_\varphi(p_\phi, p_T) < \inf_{\lambda, \mu} \hat{D}_\varphi(p_{(\lambda, \infty, \mu)}, p_T) \quad \text{if } \gamma < 0, \quad (32)$$

we have the following conclusion (see [19])

Conclusion 3. Using Propositions 1 and 4, under condition (31, 32) the sequence $(\hat{D}_\varphi(p_{\phi^k}, p_T))_k$ defined through formula (6) converges and there exists a subsequence $(\phi^{N(k)})$ which converges to a stationary point of the estimated divergence. Moreover, every limit point of the sequence $(\phi^k)_k$ is a stationary point of the estimated divergence.

In the case of the likelihood $\varphi(t) = -\log(t) + t - 1$, then if we initialize any of the proximal-point algorithms (16) or (18,19) with ϕ^0 verifying:

$$J(\phi^0) > \max \left[J \left(0, \infty, \frac{1}{n} \sum_{i=1}^n y_i \right), J \left(1, \frac{1}{n} \sum_{i=1}^n y_i, \infty \right) \right] \quad (33)$$

then we reach the following conclusion (see[1])

Conclusion 4. Using Propositions 1 and 4, under condition (33) the sequence $(J(\phi^k))_k$ converges and there exists a subsequence $(\phi^{N(k)})$ which converges to a stationary point of the likelihood function. Moreover, every limit point of the sequence $(\phi^k)_k$ is a stationary point of the likelihood.

Assumption A3 is not fulfilled (this part applies for all aforementioned situations). We study the equation $D_\psi(\phi|\phi') = 0$ for $\phi \neq \phi'$. By definition of D_ψ , it is given by a sum of nonnegative terms, which implies that all terms need to be equal to zero. The following lines are equivalent $\forall i \in \{1, \dots, n\}$:

$$\begin{aligned} h_i(0|\lambda, \mu_1, \mu_2) &= h_i(0|\lambda', \mu'_1, \mu'_2); \\ \frac{\lambda e^{-\frac{1}{2}(y_i - \mu_1)^2}}{\lambda e^{-\frac{1}{2}(y_i - \mu_1)^2} + (1 - \lambda)e^{-\frac{1}{2}(y_i - \mu_2)^2}} &= \frac{\lambda' e^{-\frac{1}{2}(y_i - \mu'_1)^2}}{\lambda' e^{-\frac{1}{2}(y_i - \mu'_1)^2} + (1 - \lambda')e^{-\frac{1}{2}(y_i - \mu'_2)^2}}; \\ \log \left(\frac{1 - \lambda}{\lambda} \right) - \frac{1}{2}(y_i - \mu_2)^2 + \frac{1}{2}(y_i - \mu_1)^2 &= \log \left(\frac{1 - \lambda'}{\lambda'} \right) - \frac{1}{2}(y_i - \mu'_2)^2 + \frac{1}{2}(y_i - \mu'_1)^2. \end{aligned}$$

Looking at this set of n equations as an equality of two polynomials on y of degree 1 at n points¹³, we deduce that as we dispose of two distinct observations, say, y_1 and y_2 , the two polynomials need to have the same coefficients. Thus the set of n equations is equivalent to the following two equations:

$$\begin{cases} \mu_1 - \mu_2 &= \mu'_1 - \mu'_2 \\ \log\left(\frac{1-\lambda}{\lambda}\right) + \frac{1}{2}\mu_1^2 - \frac{1}{2}\mu_2^2 &= \log\left(\frac{1-\lambda'}{\lambda'}\right) + \frac{1}{2}\mu_1'^2 - \frac{1}{2}\mu_2'^2 \end{cases} \quad (34)$$

These two equations with three variables have an infinite number of solutions. Take for example $\mu_1 = 0, \mu_2 = 1, \lambda = \frac{2}{3}, \mu'_1 = \frac{1}{2}, \mu'_2 = \frac{3}{2}, \lambda' = \frac{1}{2}$. This entails that, for any $\phi \in \Phi$, there exists an infinite number of elements in Φ for which the value of D_ψ between ϕ and any one of them is equal to zero. This proves that assumption A3 is not fulfilled in the Gaussian mixture.

Remark 5. The previous conclusion can be extended to any two-component mixture of exponential families having the form:

$$p_\phi(y) = \lambda e^{\sum_{i=1}^{m_1} \theta_{1,i} y^i - F(\theta_1)} + (1 - \lambda) e^{\sum_{i=1}^{m_2} \theta_{2,i} y^i - F(\theta_2)}.$$

One may write the corresponding n equations. The polynomial of y_i has a degree of at most $\max(m_1, m_2)$. Thus, if one disposes of $\max(m_1, m_2) + 1$ distinct observations, the two polynomials will have the same set of coefficients. Finally, if $(\theta_1, \theta_2) \in \mathbb{R}^{d-1}$ with $d > \max(m_1, m_2)$, then assumption A3 does not hold.

This conclusion holds for both algorithms (16) or (18,19). Unfortunately, we have no information about the difference between consecutive terms $\|\phi^{k+1} - \phi^k\|$ except for the case of $\psi(t) = \varphi(t) = -\log(t) + t - 1$ which corresponds to the classical EM recurrence:

$$\lambda^{k+1} = \frac{1}{n} \sum_{i=1}^n h_i(0|\phi^k), \quad \mu_1^{k+1} = \frac{\sum_{i=1}^n y_i h_i(0|\phi^k)}{\sum_{i=1}^n h_i(0|\phi^k)}, \quad \mu_1^{k+1} = \frac{\sum_{i=1}^n y_i h_i(1|\phi^k)}{\sum_{i=1}^n h_i(1|\phi^k)}.$$

In such case, [4] has shown that we can prove directly that $\phi^{k+1} - \phi^k$ converges to 0 without the use of Proposition 3.

B. Two-component Weibull mixture

Let p_ϕ be a two-component Weibull mixture:

$$p_\phi(x) = 2\lambda\phi_1(2x)^{\phi_1-1}e^{-(2x)^{\phi_1}} + (1-\lambda)\frac{\phi_2}{2}\left(\frac{x}{2}\right)^{\phi_2-1}e^{-\left(\frac{x}{2}\right)^{\phi_2}}, \quad \phi = (\lambda, \phi_1, \phi_2). \quad (35)$$

We have $\Phi = (0, 1) \times \mathbb{R}_+^* \times \mathbb{R}_+^*$. Similarly to the Gaussian example, we will study convergence properties in light of our theoretical approach. We will only be interested in power divergences defined through the Cressie-Read class of functions $\varphi = \varphi_\gamma$ given by (4).

The weight functions h_i are given by:

$$h_i(1|\phi) = \frac{2\lambda\phi_1(2x)^{\phi_1-1}e^{-(2x)^{\phi_1}}}{2\lambda\phi_1(2x)^{\phi_1-1}e^{-(2x)^{\phi_1}} + (1-\lambda)\frac{\phi_2}{2}\left(\frac{x}{2}\right)^{\phi_2-1}e^{-\left(\frac{x}{2}\right)^{\phi_2}}}, \quad h_i(2|\phi) = 1 - h_i(1|\phi).$$

It is clear the functions h_i are of class $\mathcal{C}^1(\text{int}(\Phi))$ and so does $\phi \mapsto D_\psi(\phi, \phi')$ for any $\phi' \in \Phi$.

If we use the MDPD (2), the continuity and differentiability of the estimated divergence \hat{D}_a can be

¹³The second order terms vanish from both sides of the each equation.

treated similarly to the Gaussian example. The proof of compactness of Φ^0 is also similar. We identify a condition on the initialization of the algorithm in order to make Φ^0 bounded.

$$\hat{D}_a(p_{\lambda, \phi}, p_T) < \min \left(0, \inf_{\phi_1 > 0, \lambda \in [\eta, 1-\eta]} \hat{D}_a(p_{(\lambda, \phi_1, \infty)}, p_T) \right).$$

A conclusion similar to Conclusion 1 can be stated here.

If we are using the dual estimator defined by (5), then continuity can be treated similarly to the case of the Gaussian example. Here, however, the continuity and differentiability of the optimized function $f(\alpha, \phi)$, where $\hat{D}_\varphi(p_\phi, p_T) = \sup_\alpha f(\alpha, \phi)$, are more technical. We list the following three results without any proof, because it suffices to study the integral term in the formula. Suppose, without loss of generality, that $\phi_1 < \phi_2$ and $\alpha_1 < \alpha_2$.

- 1) For $\gamma > 1$, which includes the Pearson's χ^2 case, the dual representation is *not* well defined since $\sup_\alpha f(\alpha, \phi) = \infty$;
- 2) For $\gamma \in (0, 1)$, function $f(\alpha, \phi)$ is continuous.
- 3) For $\gamma < 0$, function $f(\alpha, \phi)$ is continuous and well defined for $\phi_1 < \frac{\gamma-1}{\gamma}\alpha_1$ and $\alpha_2 \geq \phi_2$. Otherwise $f(\alpha, \phi) = -\infty$, but the supremum $\sup_\alpha f(\alpha, \phi)$ is still well defined.

In both cases 2 and 3, differentiability of function $f(\alpha, \phi)$ holds only on a subset of $\Phi \times \Phi$ which cannot be written as $A \times B$, and thus the theoretical approaches presented in Section II are not suitable. In order to end this part, we emphasize the fact that, similarly to the Gaussian example, even continuity of the estimated divergence $\hat{D}_\varphi(p_\phi, p_T)$ with respect to ϕ cannot be directly using the theoretical approaches presented in paragraph (II) unless we suppose that Φ is compact. Indeed, if Φ is compact, then using Theorem 1.17 from [20], continuity of the estimated divergence is a direct result. Differentiability of the estimated divergence is far more difficult and needs more investigations on the form of the estimated divergence and the model used.

Similar conclusion to Conclusion 2 can be stated here with no changes except for the fact that assumption AC is not fulfilled. This entails that our conclusion will be about the subgradient of the estimated divergence.

If we are using the kernel-based dual estimator given by (6) with a Gaussian kernel density estimator, then things are a lot simplified. We need only to treat the integral term. From an analytic point of view, the study of continuity depends on the kernel used; more specifically its tail behavior. If we take a Gaussian kernel, then we have:

- For $\gamma > 1$, it is necessary that $\min(\phi_1, \phi_2) > 2$, otherwise the estimated divergence is infinity. Thus, it is necessary for either of the true values of the shapes to be inferior to 2 in order for the estimation to be valid;
- For $\gamma \in (0, 1)$, then the estimated divergence is $\mathcal{C}^1(\text{int}(\Phi))$;
- For $\gamma < 0$, it is necessary that $\min(\phi_1, \phi_2) < 1 - \frac{1}{\gamma}$ and $\max(\phi_1, \phi_2) < 2$. If these conditions do not hold, then the estimated divergence is minimized at $-\infty$ at any vector of parameters which does not verify the previous condition.

In the first case, if we use a heavier-tailed kernel such as the Cauchy Kernel, the estimated divergence becomes $\mathcal{C}^1(\text{int}(\Phi))$. In the third case, if we use a compact-supported kernel such as the Epanechnikov's

kernel, the condition is reduced to only $\min(\phi_1, \phi_2) < 1 - \frac{1}{\gamma}$.

Similar conditions to (31,32) can be obtained and we have the same conclusion as Conclusion 3.

In the case of the Likelihood $\varphi(t) = -\log(t) + t - 1$, we illustrate the convergence of the EM algorithm through our theoretical approach. Assumptions A1 and AC are clearly verified since both the log-likelihood and the proximal term are sums of continuously differentiable functions, and integrals do not intervene here. The set Φ^0 is given by:

$$\begin{aligned}\Phi^0 &= \{\phi \in \Phi, J(\phi) \geq J(\phi^0)\} \\ &= J^{-1}([J(\phi^0), \infty)) \\ &= \{\phi \in \Phi, L(\phi) \geq L(\phi^0)\}\end{aligned}$$

where $L(\phi)$ is the likelihood of the model, and $J(\phi) = \log(L(\phi))$ is the log-likelihood function. We will show that under similar conditions to the Gaussian mixture, the set Φ^0 is compact.

Closedness of Φ^0 . Since the shape parameter is supposed to be positive, continuity of the log-likelihood would imply only that Φ^0 is closed in $[0, 1] \times \mathbb{R}_+^* \times \mathbb{R}_+^*$, a space which is not closed and hence is not complete. We therefore, propose to extend the definition of shape parameter on 0. From a statistical point of view, this extension is not reasonable since the density function of Weibull distribution with a shape parameter equal to 0 is the zero function which is not a probability density. Besides, identifiability problems would appear for a mixture model. Nevertheless, our need is only for analytical purpose. We will add suitable conditions on ϕ^0 in order to avoid such subtlety keeping in hand the closedness property. We suppose now that the shape parameter can have values in \mathbb{R}_+ . The set Φ^0 is now the inverse image of $[L(\phi^0), \infty)$ by the likelihood function¹⁴ which is continuous on $[0, 1] \times \mathbb{R}_+ \times \mathbb{R}_+$. Hence, it is closed in the space $[0, 1] \times \mathbb{R}_+ \times \mathbb{R}_+$ provided the euclidean norm which is complete. It suffices then to prove that Φ^0 is bounded.

Boundedness of Φ^0 . We will make similar arguments to the case of the Gaussian mixture example. We need to calculate the limit at infinity when the shape parameter of either of the two components tends to infinity. If both ϕ_1 and ϕ_2 goes to infinity, the log-likelihood tends to $-\infty$. Hence any choice of a finite ϕ^0 can avoid this case. Suppose now that ϕ_1 goes to infinity whereas ϕ_2 stays bounded. The corresponding limit of the log-likelihood functions is given by:

$$J(\lambda, \infty, \phi_2) = \sum_{i=1}^n \log \left((1 - \lambda) \frac{\phi_2}{2} \left(\frac{y_i}{2} \right)^{\phi_2 - 1} e^{-\left(\frac{y_i}{2}\right)^{\phi_2}} \right)$$

if there is no observation y_i equal to $\frac{1}{2}$. In fact, if there is $y_i = \frac{1}{2}$, the limit is $+\infty$ and the set Φ^0 cannot be bounded. However, it is improbable to get such an observation since the probability of getting an observation equal to $\frac{1}{2}$ is zero. The case when ϕ_2 goes to infinity whereas ϕ_1 stays bounded is treated similarly.

To avoid the two previous scenarios, one should choose the initial point of the algorithm ϕ^0 in a way that it verifies:

$$J(\phi^0) > \max \left(\sup_{\lambda, \phi_2} J(\lambda, \infty, \phi_2), \sup_{\lambda, \phi_1} J(\lambda, \phi_1, \infty) \right). \quad (36)$$

¹⁴We do not use this time the log-likelihood function since it is not defined when both shape parameters are zero.

Since all vectors of Φ^0 have a log-likelihood value greater than $J(\phi^0)$, the previous choice permits the set Φ^0 to avoid non-finite values of ϕ . Thus it becomes bounded whenever ϕ_0 is chosen according to condition (36). Finally, the calculus of both terms $\sup_{\lambda, \phi_1} J(\lambda, \phi_1, \infty)$ and $\sup_{\lambda, \phi_2} J(\lambda, \phi_2, \infty)$ is not feasible but numerically. They, however, can be simplified a little. One can notice by writing these terms without the logarithm (as a product), the term which has λ is maximized when it is equal to 1. The remaining of the calculus is a maximization of the likelihood function of a Weibull model¹⁵.

We conclude that the set Φ^0 is compact under condition (36). Finally, it is important to notice that condition (36) permits also to avoid the border values which corresponds to $\phi_1 = 0$ or $\phi_2 = 0$. Indeed, when either of the shape parameters is zero, the corresponding component vanishes and the corresponding log-likelihood value is less than the upper bound in condition (36). The same conclusion as Conclusion 4 can be stated here for the Weibull mixture model.

Notice that the verification of assumption A3 is a hard task here because it results in a set of n nonlinear equations in y_i and cannot be treated in a similar way to the Gaussian mixture.

C. Pearson's χ^2 algorithm for a Cauchy model

Let $\{(x_i, y_i), i = 0, \dots, n\}$ be an n -sample drawn from the joint probability law defined by the density function:

$$f(x, y|a, x_0) = \frac{a(y - x_0)^2 e^x}{\pi(a^2 + (y - x_0)^2 e^x)^2}, \quad x \in [0, \infty), y \in \mathbb{R}$$

where $a \in [\varepsilon, \infty)$, with $\varepsilon > 0$, denotes a scale parameter and $x_0 \in \mathbb{R}$ denotes a location parameter. We define an exponential probability law with parameter $\frac{1}{2}$ on the labels. It is given by the density function:

$$q(x) = \frac{1}{2} e^{-x/2}.$$

Now, the model defined on the observed data becomes a Cauchy model with two parameters:

$$p_{(a, x_0)}(y) = \int_0^\infty f(x, y|a, x_0) dx = \frac{a}{\pi(a^2 + (y - x_0)^2)}, \quad a \geq \varepsilon > 0, x_0 \in \mathbb{R}.$$

The goal of this example is to show how we prove assumptions A1-3 and AC in order to explore the convergence properties of the sequence ϕ^k generated by either of the algorithms (16) and (18,19). We also discuss the analytical properties of the dual representation of the divergence.

In this example, we only focus on the dual representation of the divergence given by (5) because the resulting MD φ DE is robust against outliers (so does the MLE). Thus there is no need to use a robust estimator such as the kernel-based MD φ DE which needs a choice of a suitable kernel and window.

1) *Cauchy model with zero location:* We suppose here that $x_0 = 0$, and we are only interested in estimating the scale parameter a . The Pearson's χ^2 divergence is given by:

$$D(p_a, p_{a^*}) = \frac{1}{2} \int \left[\frac{p_a(y)}{p_{a^*}} - 1 \right]^2 p_{a^*}(y) dy.$$

Let's rewrite the dual representation of the Chi square divergence:

$$\hat{D}(p_a, p_{a^*}) = \sup_{b \geq \varepsilon} \left\{ \int_{\mathbb{R}} \frac{p_b^2(x)}{p_a(x)} dx - \frac{1}{2n} \sum_{i=1}^n \frac{p_b^2(y_i)}{p_a^2(y_i)} \right\} - \frac{1}{2}.$$

¹⁵In a Weibull model, the calculus of the MLE cannot be done but numerically when the parameter of interest is the shape parameter.

A simple calculus shows:

$$\int_{\mathbb{R}} \frac{p_b^2(x)}{p_a(x)} dx = \frac{(a^2 + b^2)\pi}{2ab}.$$

This implies a simpler form for the dual representation of the divergence:

$$\hat{D}(p_a, p_{a^*}) = \sup_{b \geq \varepsilon} \left\{ \frac{(a^2 + b^2)}{2ab} - \frac{1}{2n} \sum_{i=1}^n \frac{a^2(b^2 + y_i^2)^2}{b^2(a^2 + y_i^2)^2} \right\} - \frac{1}{2}. \quad (37)$$

Let $f(a, b)$ denote the optimized function in the above formula. We calculate the first derivative with respect to b :

$$\frac{\partial f}{\partial b}(a, b) = -\frac{\pi a}{2b^2} + \frac{\pi}{2a} - \frac{1}{2n} \sum_{i=1}^n \frac{a^2}{(a^2 + y_i^2)^2} \left(2b - \frac{2y_i^4}{b^3} \right).$$

Notice that as $a \geq \varepsilon$ the term $\frac{\pi}{2a}$ stays bounded away from infinity uniformly. Therefore, it suffices then that b exceeds a finite value b_0 in order that the derivative becomes negative. Hence, there exists b_0 such that $b \mapsto f(a, b)$ becomes decreasing independently of a . On the other hand $\forall a > 0, \lim_{b \rightarrow \infty} f(a, b) = -\infty$. It results that all values of the function $b \mapsto f(a, b)$ for $b > b_0$ does not have any use in the calculus of the supremum in (37), since, by the decreasing property if $b \mapsto f(a, b)$, they all should have values less than the value at b_0 . We may now rewrite the dual representation of the Chi square divergence as :

$$\hat{D}(p_a, p_{a^*}) = \sup_{b \in [\varepsilon, b_0]} \left\{ \frac{(a^2 + b^2)}{2ab} - \frac{1}{2n} \sum_{i=1}^n \frac{a^2(b^2 + y_i^2)^2}{b^2(a^2 + y_i^2)^2} \right\} - \frac{1}{2}. \quad (38)$$

We have now two pieces of information about $f(a, b)$. The first is that it is level-bounded locally in b uniformly in a (see paragraph (II-B)). The second is that we are exactly in the context of lower- \mathcal{C}^1 functions (II-A). First of all, function f is $\mathcal{C}^1([\varepsilon, \infty) \times [\varepsilon, \infty))$ function, so that part (a) of Theorem 3 is verified and the function $a \mapsto \hat{D}(p_a, p_{a^*})$ is strictly continuous. To prove it is continuously differentiable, we need to prove that the set

$$Y(a) = \bigcup_{b \in \arg \max_{b'} f(a, b)} \left\{ \frac{\partial f}{\partial a}(a, b) \right\}$$

contains but one element. From a theoretic point of view, two possible methods are available: Prove that either there is a unique maximum for a fixed a , or that the derivative with respect to a at all maxima *does not depend* on a (they have the same value). In our example, function $b \mapsto f(a, b)$ is not concave. We may also plot it using any mathematical tool provided that we already have the data set. We tried out a simple example and generated a 10-sample of the standard Cauchy distribution ($a = 1$), see table (I). We used Mathematica to draw a 3D figure of function f , see figure (1).

y_i	0.534	-18.197	0.726	-0.439	-1.945	0.0119	12.376	-0.953	0.698	0.818
-------	-------	---------	-------	--------	--------	--------	--------	--------	-------	-------

Table I

A 10-SAMPLE CAUCHY DATASET.

It is clear that for a fixed a , the function $b \mapsto f(a, b)$ has two maxima which may both be global maxima. For example for $a = 0.9$, one gets figure (2). It is clearer now that conditions of Theorem 3 are not fulfilled, and we cannot prove that function $\hat{D}(p_a, p_{a^*})$ is continuously differentiable every where.

It is however not the end of the road. We still have the results presented in paragraph (II-A). Function

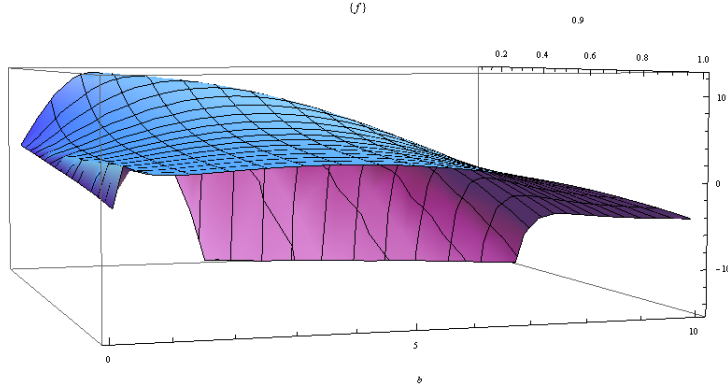


Figure 1. A 3D plot of function $f(a, b)$ for a 10-sample of the standard Cauchy distribution.



Figure 2. A 2D plot of function $f(0.9, b)$ for a 10-sample of the standard Cauchy distribution.

$\hat{D}(p_a, p_{a^*})$ is lower- \mathcal{C}^1 . Therefore, it is strictly continuous and almost everywhere continuously differentiable. Hence, we may hope that the limit points of the sequence $(\phi^k)_k$ for algorithm (16) are in the set of points where the dual representation of the Chi square divergence is \mathcal{C}^1 , or be more reasonable and state any further result on the sequence in terms of the subgradient of $\hat{D}(p_a, p_{a^*})$.

a) Compactness of Φ^0 . We check when the set $\Phi^0 = \{a | \hat{D}(p_a, p_{a^*}) \leq \hat{D}(p_{a_0}, p_{a^*})\}$ is closed and bounded in $[\varepsilon, \infty)$ for an initial point a_0 . **Closedness** is proved using continuity of $\hat{D}(p_a, p_{a^*})$. Indeed,

$$\Phi^0 = \hat{D}^{-1}(p_a, p_{a^*}) \left((-\infty, \hat{D}(p_{a_0}, p_{a^*})] \right).$$

Boundedness is proved by contradiction. Suppose that Φ^0 is unbounded, then there exists a sequence $(a^l)_l$ of points of Φ^0 which goes to infinity. Formula (38) shows that b stays in a bounded set during the calculus of the supremum. Hence the continuity of $\hat{D}(p_a, p_{a^*})$ implies:

$$\lim_{a \rightarrow \infty} \hat{D}(p_a, p_{a^*}) = +\infty.$$

This shows that by choosing any finite a_0 , the set Φ^0 becomes bounded. Indeed, the relation defining Φ^0 implies that $\forall l, \hat{D}(p_{a^l}, p_{a^*}) \leq \hat{D}(p_{a_0}, p_{a^*}) < \infty$, and a contradiction is reached by taking the limit of each part of this inequality. Hence Φ^0 is closed and bounded in the space $[\varepsilon, \infty)$ which is complete

provided with the euclidean distance. We conclude that Φ^0 is compact¹⁶.

In this simple example, we only can use algorithm (16) since there is only one parameter of interest. Proposition 4 can be used to deduce convergence of any convergent subsequence to a generalized stationary point of $\hat{D}(p_a, p_{a^*})$.

To deduce more results about the sequence $(a^k)_k$, we may try and verify assumption A3 using Lemma 1. Let's write functions h_i .

$$h_i(x|a) = \frac{f(x, y_i|a)}{p_a(y_i)} = \frac{y_i^2 e^x (a^2 + y_i^2)}{(a^2 + e^x y_i^2)^2}.$$

Clearly, for any $i \in \{1, \dots, n\}$ and $a \geq \varepsilon$, function $x \mapsto h_i(x|a)$ is continuous. Let $a, b \geq \varepsilon$ such that $a \neq b$. Suppose that:

$$\forall i, \quad h_i(x|a) = h_i(x|b) \quad \forall x \geq 0.$$

This entails that:

$$a^2 b^4 - a^4 b^2 + (b^4 - a^4) y_i^2 + (a^2 e^{2x} + 2b^2 e^x - b^2 e^{2x} - 2a^2 e^x) y_i^4 = 0, \quad i = 1, \dots, n.$$

This is a polynomial on y_i of degree 4 which coincides with the zero polynomial on n points. If there exists 5 distinct observations¹⁷, then the two polynomials will have the same coefficients. Hence, we have $b^4 - a^4 = 0$. This implies that $a = b$ since they are both positive real numbers. We conclude that $D_\psi(a, b) = 0$ whenever $a = b$ which is equivalent to assumption A3. Proposition 3 can now be applied to deduce that sequence (a^k) defined by (16) (with ϕ^k replaced by a^k) is well defined and bounded. Furthermore, it verifies $a^{k+1} - a^k \rightarrow 0$, and the limit of any convergent subsequence is a generalized stationary point of $\hat{D}(p_a, p_{a^*})$. The existence of such subsequence is guaranteed by the compactness of Φ^0 and the fact that $\forall k, a^k \in \Phi^0$.

V. EXPERIMENTAL RESULTS

We summarize the results of 100 experiments on 100-samples (with and without outliers) from two-components Gaussian and Weibull mixtures by giving the average of the error committed with the corresponding standard deviation. The criterion error is mainly the total variation distance (TVD) which is calculated using the $L1$ distance by the Scheffé lemma (see for example [27] page 129.).

$$\begin{aligned} \text{TVD}(p_\phi, p_T) &= \sup_{a < b} |dP_\phi([a, b]) - dP_T([a, b])| \\ &= \frac{1}{2} \int |p_\phi(x) - p_T(x)| dx. \end{aligned}$$

We also provide for the Gaussian mixture the values of the (squared root of the) χ^2 divergence between the estimated model and the true mixture, since it gave infinite values for the Weibull experiment. The χ^2 criterion is defined by:

$$\chi^2(p_\phi, p_T) = \int \frac{(p_\phi(x) - p_T(x))^2}{p_{\phi^*}(x)} dx.$$

¹⁶If we are to use a result which concerns the differentiability of $\hat{D}(p_a, p_{a^*})$, one should consider the case when Φ^0 shares a boundary with Φ . A possible solution to avoid this is to consider an initial point a^0 such that $\hat{D}(p_\varepsilon, p_{a^*}) > \hat{D}(p_{a_0}, p_{a^*})$. This expels the the boundary from the possible values of Φ^0 .

¹⁷If one uses the point $x = 0$, the result follows directly without supposing the existence of distinct observations.

The use of a distance such as the χ^2 divergence is due to its relative-error property. In other words, it calculates the error at a point relatively to its true value. Hence, errors at small values of the true density have their share in the overall error and are no longer negligible to points with higher density value. The total variation indicates the maximum error we might commit when calculating probabilities by replacing the true distribution by the estimated one.

We used different φ -divergences to estimate the parameters and compared the performances of the two methods of estimating a φ -divergence presented in this paper. For the Gaussian mixture, we used the Pearson's χ^2 and the Hellinger divergences, whereas in the Weibull mixture, we used the Neymann's χ^2 and the Hellinger divergences. For the MDPD, we used $a = 0.5$; a choice which gave the best tradeoff between robustness and efficiency in the simulation results in [17]. We illustrate also the performance of the EM method in the light of our method, i.e. using initializations verifying conditions (33) for the Gaussian mixture and conditions (36) for the Weibull one. When outliers were added, these initializations did not always result in good results and the convergence of the proportion was towards the border $\eta = 0.1$ or $1 - \eta = 0.9$. In such situations, the EM algorithm was initialized using another starting point manually. Last but not least, for the proximal term, we used $\psi(t) = \frac{1}{2}(\sqrt{t} - 1)^2$.

We used the Nelder-Mead algorithm (see [28]) for all optimization calculus. The method proved to be more efficient in our context than other optimization algorithms although it has a slow convergence speed. Such method is derivative-free and applies even if the the objective function is not differentiable which may be the case of the estimated divergence defined through (5). The Nelder-Mead algorithm is known to give good results in problems with dimension at least 2 and does not perform well in dimension 1. We thus used Brent's method for the unidimensional optimizations. It is also a derivative-free method which works in a compact subset from \mathbb{R} only. The calculus was done under the statistical tool [29].

Numerical integrations were performed using the `distrExIntegrate` function of package `distrEx` in the Gaussian mixture. It is a slight modification of the standard `integrate` function in the R statistical tool which performs a Gauss-Legendre quadrature approximation whenever function `integrate` fails to converge. For the Weibull mixture, the previous function did not converge always, and function `integral` of package `pracma` was used. Although being very slow, it performs very well especially on unbounded integrations and "extremely bad-behavior" integrands.

A. The two-component Gaussian mixture revisited

We consider the Gaussian mixture (29) presented earlier with true parameters $\lambda = 0.35, \mu_1 = 2, \mu_2 = 1.5$ and fixed variances $\sigma_1 = \sigma_2 = 1$. Since we are using a function error criterion, label-switching problems do not interfere. Figure (3) shows the values of the estimated divergence for both formulas (5) and (6) on a logarithmic scale at each iteration of the algorithms (16) and (18, 19) until convergence. The 1-step algorithm refers to algorithm (16), whereas 2-step refers to algorithm (18,19). We omitted the initial point in order to produce a clear image of the decrease of the objective function. For the kernel-based dual formula, we used a Gaussian kernel with window calculate using Silverman's rule of thumb. Results are presented in table (II).

Contamination was done by adding in the original sample to the 5 lowest values random observations

from the uniform distribution $\mathcal{U}[-5, -2]$. We also added to the 5 largest values random observations from the uniform distribution $\mathcal{U}[2, 5]$. Results are presented in table (III).

It is clear that the both the MDPD and the kernel-based MD φ DE are more robust than the EM algorithm and the classical MD φ DE for both the Pearson's χ^2 and the Hellinger divergences. Differences between the two choices of φ -divergences (χ^2 and Hellinger) were not significant for the two φ -divergence-based estimators.

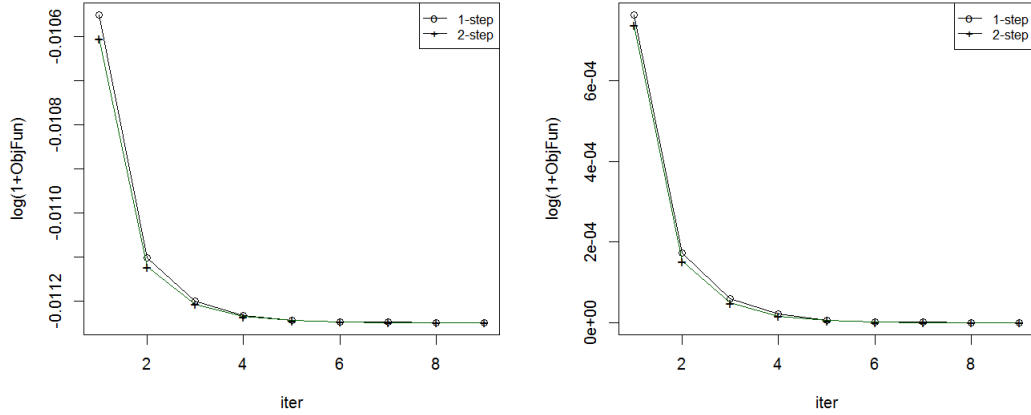


Figure 3. Decrease of the (estimated) Hellinger divergence between the true density and the estimated model at each iteration in the Gaussian mixture. The figure to the left is the curve of the values of the kernel-based dual formula (6). The figure to the right is the curve of values of the classical dual formula (5). Values are taken at a logarithmic scale $\log(1+x)$.

B. The two-component Weibull mixture model revisited

We consider the Weibull mixture (35) with $\phi_1 = 0.5, \phi_2 = 3$ and $\lambda = 0.35$ which are supposed to be unknown during the estimation procedure. We denote $\phi = (\phi_1, \phi_2)$ ($\alpha = (\alpha_1, \alpha_2)$, respectively) the shapes of the Weibull mixture model $p_{(\lambda, \phi)}$ ($p_{(\lambda, \alpha)}$, respectively). Contamination was done by replacing 10 observations of each sample chosen randomly by 10 i.i.d. observations drawn from a Weibull distribution with shape $\nu = 0.9$ and scale $\sigma = 3$. Results are presented in tables (IV) and (V).

Manipulating the optimization procedure for the Neymann's χ^2 was difficult because of the numerical integration calculus and the fact that for a subset of Φ (or $\Phi \times \Phi$ according to whether we use the estimator (5) or the estimator (6)) the integral term produces infinity, see paragraph IV-B for more details. We therefore needed to keep the optimization from approaching the border in order to avoid numerical problems. For the Hellinger divergence, there is no particular remark.

For the case of the estimated divergence (5), if $\gamma = -1$, i.e. the Neymann χ^2 , we need that $\alpha_1 < \phi_1/2$, otherwise the integral term is equal to infinity. In order to avoid numerical complications, we optimized over $\alpha_1 \leq 0.05 + \phi_1/2$. The value 0.05 ensures a small deviation from the border.

For the case of the estimated divergence (6), we used a Gaussian kernel for the Hellinger divergence. For the Neymann's χ^2 divergence, we used the Epanechnikov's kernel to avoid problems at infinity. Besides,

Table II

THE MEAN VALUE OF ERRORS COMMITTED IN A 100-RUN EXPERIMENT WITH THE STANDARD DEVIATION. NO OUTLIERS ARE CONSIDERED HERE. THE DIVERGENCE CRITERION IS THE CHI SQUARE DIVERGENCE OR THE HELLINGER. THE PROXIMAL TERM IS CALCULATED WITH $\psi(t) = \frac{1}{2}(\sqrt{t} - 1)^2$.

Estimation method		Error criterion	
		$\sqrt{\chi^2}$	TVD
Chi square			
Algorithm (16)	MD φ DE	0.108, sd = 0.052	0.061, sd = 0.029
	kernel-based MD φ DE	0.118 , sd = 0.052	0.066 ,sd= 0.027
Algorithm (18,19)	MD φ DE	0.108, sd = 0.052	0.061, sd = 0.029
	kernel-based MD φ DE	0.118, sd = 0.051	0.066 ,sd= 0.027
Hellinger			
Algorithm (16)	MD φ DE	0.108, sd = 0.052	0.050 , sd=0.025
	kernel-based MD φ DE	0.113, sd = 0.044	0.064 ,sd=0.025
Algorithm (18,19)	MD φ DE	0.108, sd = 0.052	0.061, sd = 0.029
	kernel-based MD φ DE	0.113, sd = 0.045	0.064 ,sd=0.025
MDPD $a = 0.5$ - Algorithm (16)		0.117, sd = 0.049	0.065, sd = 0.025
MDPD $a = 0.5$ - Algorithm (18,19)		0.117, sd = 0.047	0.065, sd = 0.025
EM		0.113, sd = 0.044	0.064 , sd = 0.025

it permits to integrate only over $[0, \max(Y) + w]$, where w is the window of the kernel, instead of $[0, \infty)$.

In order to avoid problems near zero, it is necessary that $\min(\phi_1, \phi_2) < 1 - \frac{1}{\gamma} = 2$.

Experimental results show a clear robustness of the estimators calculated using the density power divergece (the MDPD) and the kernel-based MD φ DE in comparison to other estimators using the Hellinger divergence. When we are under the model, all estimation methods have the same performance. On the other hand, using the Neymann χ^2 divergence, results are different in the presence of outliers. The classical MD φ DE calculated using formula (5) shows better robustness than other estimators except for the MDPD, but is still not as good as the robustness of the kernel-based MD φ DE using the Hellinger or the MDPD. Lack of robustness of the kernel-based MD φ DE is not very surprising since the influence function of the kernel-based MD φ DE is unbounded when we use the Neymann χ^2 divergence in simple models such as the Gaussian model, see Example 2 in [17].

In what concerns the proximal algorithm, there is no significant difference between the results obtained using the 1-step algorithm (16) and the ones obtained using the 2-step algorithm (18,19) using the Hellinger divergence. Differences appear when we used the Neymann χ^2 divergence with the classical MD φ DE. This shows again the difficulty in handling the supermal form of the dual formal (5).

Table III

ERROR COMMITTED IN ESTIMATING THE PARAMETERS OF A 2-COMPONENT GAUSSIAN MIXTURE WITH 10% OUTLIERS. THE DIVERGENCE CRITERION IS THE CHI SQUARE DIVERGENCE OR THE HELLINGER. THE PROXIMAL TERM IS CALCULATED WITH

$$\psi(t) = \frac{1}{2}(\sqrt{t} - 1)^2.$$

Estimation method		Error criterion	
		χ^2	TVD
Chi square			
Algorithm (16)	MD φ DE	0.334, sd = 0.097	0.146, sd=0.036
	kernel-based MD φ DE	0.149 , sd = 0.059	0.084 ,sd=0.033
Algorithm (18,19)	MD φ DE	0.333, sd = 0.097	0.149, sd = 0.033
	kernel-based MD φ DE	0.149 , sd = 0.059	0.084, sd=0.033
Hellinger			
Algorithm (16)	MD φ DE	0.321, sd = 0.096	0.146, sd=0.034
	kernel-based MD φ DE	0.155 , sd = 0.059	0.087 ,sd=0.033
Algorithm (18,19)	MD φ DE	0.322, sd = 0.097	0.147, sd = 0.034
	kernel-based MD φ DE	0.156 , sd = 0.059	0.087 ,sd=0.033
MDPD $a = 0.5$ - Algorithm (16)		0.129, sd = 0.049	0.065, sd = 0.025
MDPD $a = 0.5$ - Algorithm (18,19)		0.138, sd = 0.053	0.078, sd = 0.030
EM		0.335, sd = 0.102	0.150, sd = 0.034

Table IV

THE MEAN VALUE OF ERRORS COMMITTED IN A 100-RUN EXPERIMENT OF A TWO-COMPONENT WEIBULL MIXTURE WITH THE STANDARD DEVIATION. NO OUTLIERS ARE CONSIDERED. THE DIVERGENCE CRITERION IS THE NEYMANN'S χ^2 DIVERGENCE OR

THE HELLINGER. THE PROXIMAL TERM IS TAKEN WITH $\psi(t) = \frac{1}{2}(\sqrt{t} - 1)^2$.

Estimation method		Error criterion
		TVD
Neymann Chi square		
Algorithm (16)	MD φ DE	0.114 , sd = 0.032
	kernel-based MD φ DE	0.057, sd = 0.028
Algorithm (18,19)	MD φ DE	0.131, sd = 0.042
	kernel-based MD φ DE	0.056, sd = 0.026
Hellinger		
Algorithm (16)	MD φ DE	0.059, sd = 0.024
	kernel-based MD φ DE	0.057, sd = 0.029
Algorithm (18,19)	MD φ DE	0.061, sd = 0.026
	kernel-based MD φ DE	0.057, sd = 0.029
MDPD $a = 0.5$ - Algorithm (16)		0.056, sd = 0.029
MDPD $a = 0.5$ - Algorithm (18,19)		0.056, sd = 0.029
EM		0.059, sd = 0.024

Table V

THE MEAN VALUE OF ERRORS COMMITTED IN A 100-RUN EXPERIMENT OF A TWO-COMPONENT WEIBULL MIXTURE WITH THE STANDARD DEVIATION. 10% OUTLIERS ARE CONSIDERED. THE DIVERGENCE CRITERION IS THE NEYMANN'S χ^2 DIVERGENCE OR THE HELLINGER. THE PROXIMAL TERM IS TAKEN WITH $\psi(t) = \frac{1}{2}(\sqrt{t} - 1)^2$.

Estimation method		Error criterion
		TVD
Neymann Chi square		
Algorithm (16)	MD φ DE	0.085, sd = 0.036
	kernel-based MD φ DE	0.138, sd = 0.066
Algorithm (18,19)	MD φ DE	0.096, sd = 0.057
	kernel-based MD φ DE	0.127, sd = 0.056
Hellinger		
Algorithm (16)	MD φ DE	0.120, sd = 0.034
	kernel-based MD φ DE	0.068, sd = 0.034
Algorithm (18,19)	MD φ DE	0.121, sd = 0.034
	kernel-based MD φ DE	0.068, sd = 0.034
MDPD $a = 0.5$ - Algorithm (16)		0.060, sd = 0.029
MDPD $a = 0.5$ - Algorithm (18,19)		0.061, sd = 0.029
EM		0.129, sd = 0.046

VI. CONCLUSIONS

We presented in this paper a proximal-point algorithm whose objective was the minimization of (an estimate of) a φ -divergence. The set of algorithms proposed here covers the EM algorithm. We provided in several examples a proof of convergence of the EM algorithm in the spirit of our approach. We also showed how we may prove convergence for the two estimates of the φ -divergence (5) and (6) and for the density power divergence (1). We reestablished similar results to the ones in [4] in the context of general divergences, and provided a new result by relaxing the identifiability condition on the proximal term. Our simulation results permit to conclude that the proximal algorithm works. The two-step algorithm (18,19) showed in the most difficult situations considered here a slight deterioration in performance comparing to the original one (16) which is very encouraging especially that the dimension of the optimization is reduced at each step. Simulations have shown again the robustness of φ -divergences and the density power divergence against outliers in comparison to the MLE. The algorithm could be used to calculate other divergence-based estimators such as [16] and [14] or Rényi pseudodistances ([25]). The role of the proximal term and its influence on the convergence of the algorithm were not discussed here and will be considered in a future work.

REFERENCES

- [1] D. Al Mohamad and M. Broniatowski, *Generalized EM Algorithms for Minimum Divergence Estimation*. Cham: Springer International Publishing, 2015, pp. 417–426. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-25040-3_45
- [2] G. McLachlan and T. Krishnan, *The EM Algorithm and Extensions*, ser. Wiley Series in Probability and Statistics. Wiley, 2007.
- [3] S. Chrétien and A. O. Hero, “Generalized proximal point algorithms and bundle implementations,” Department of Electrical Engineering and Computer Science, The University of Michigan, Tech. Rep., 1998.
- [4] P. Tseng, “An analysis of the EM algorithm and entropy-like proximal point methods,” *Math. Oper. Res.*, vol. 29, no. 1, pp. 27–44, 2004.
- [5] A. Goldstein and I. Russak, “How good are the proximal point algorithms?” *Numerical Functional Analysis and Optimization*, vol. 9, no. 7-8, pp. 709–724, 1987.
- [6] S. Chretien and A. Hero, “Acceleration of the em algorithm via proximal point iterations,” in *Information Theory, 1998. Proceedings. 1998 IEEE International Symposium on*, 1998, pp. 444–.
- [7] S. Chrétien and A. O. Hero, “On em algorithms and their proximal generalizations,” *ESAIM: Probability and Statistics*, vol. 12, pp. 308–326, 1 2008.
- [8] A. Basu, I. R. Harris, N. L. Hjort, and M. C. Jones, “Robust and efficient estimation by minimizing a density power divergence,” *Biometrika*, vol. 85, no. 3, pp. 549–559, 09 1998.
- [9] M. Broniatowski and I. Vajda, “Several applications of divergence criteria in continuous families,” *Kybernetika*, vol. 48, no. 4, pp. 600–636, 2012.
- [10] I. Csiszár, “Eine informationstheoretische Ungleichung und ihre anwendung auf den Beweis der ergodizität von Markoffschen Ketten,” *Publications of the Mathematical Institute of Hungarian Academy of Sciences*, vol. 8, pp. 95–108, 1963.
- [11] M. Broniatowski and A. Keziou, “Parametric estimation and tests through divergences and the duality technique,” *J. Multivariate Anal.*, vol. 100, no. 1, pp. 16–36, 2009.
- [12] —, “Minimization of divergences on sets of signed measures,” *Studia Sci. Math. Hungar.*, vol. 43, no. 4, pp. 403–442, 2006.
- [13] F. Liese and I. Vajda, “On divergences and informations in statistics and information theory,” *IEEE Transactions on Information Theory*, vol. 52, no. 10, pp. 4394–4412, 2006.

- [14] R. Beran, “Minimum hellinger distance estimates for parametric models,” *Ann. Statist.*, vol. 5, no. 3, pp. 445–463, 05 1977.
- [15] C. Park and A. Basu, “Minimum disparity estimation : Asymptotic normality and breakdown point results,” *Bulletin of informatics and cybernetics*, vol. 36, pp. 19–33, 2004.
- [16] A. Basu and B. G. Lindsay, “Minimum disparity estimation for continuous models: Efficiency, distributions and robustness,” *Annals of the Institute of Statistical Mathematics*, vol. 46, no. 4, pp. 683–705, 1994.
- [17] D. Al Mohamad, “Towards a better understanding of the dual representation of phi divergences,” *Statistical Papers*, 2016, under revision. [Online]. Available: <http://arxiv.org/abs/1506.02166>
- [18] A. Toma and M. Broniatowski, “Dual divergence estimators and tests: Robustness results.” *J. Multivariate Analysis*, vol. 102, no. 1, pp. 20–36, 2011.
- [19] D. A. Mohamad and M. Broniatowski, “A proximal point algorithm for minimum divergence estimators with application to mixture models,” *Entropy*, 2016, submitted. [Online]. Available: <https://arxiv.org/abs/1603.07117>
- [20] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*, 3rd ed., ser. Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen. Springer, 1998.
- [21] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the em algorithm,” *Journal of the Royal Statistical Society, series B*, vol. 39, no. 1, pp. 1–38, 1977.
- [22] D. Titterton, A. Smith, and U. Makov, *Statistical Analysis of Finite Mixture Distributions*. Wiley, New York, 1985.
- [23] C. F. J. Wu, “On the convergence properties of the em algorithm,” *Ann. Statist.*, vol. 11, no. 1, pp. 95–103, 03 1983.
- [24] A. Ghosh, I. R. Harris, A. Maji, A. Basu, and L. Pardo, “A generalized divergence for statistical inference,” Byesian and Interdisciplinary Research Unit Indian Statistical Institute, Tech. Rep., 2013.
- [25] A. Toma and S. Leoni-Aubin, “Optimal robust m-estimators using Rényi pseudodistances,” *Journal of Multivariate Analysis*, vol. 115, no. C, pp. 359–373, 2013.
- [26] A. Ostrowski, *Solution of equations and systems of equations*, ser. Pure and applied mathematics. Academic Press, 1966.
- [27] A. Meister, *Deconvolution Problems in Nonparametric Statistics*, ser. Lecture Notes in Statistics. Springer, 2009.
- [28] J. A. Nelder and R. Mead, “A simplex method for function minimization,” *Computer Journal*, vol. 7, pp. 308–313, 1965.
- [29] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2013. [Online]. Available: <http://www.R-project.org/>